Breaking Bad Opioid Sorting*

Han Ng† Stephan Sagl‡ Jia Xiang§

September 11, 2026

Abstract

We investigate how consumer sorting contributes to outcome differences a nd its implications for expert-targeted policies. Using employer-sponsored insurance claims on opioid prescribing for chronic pain, we show that patient sorting explains more than twice as much of the prescribing variation as physicians' prescribing propensity, and stems almost entirely from preferences exceeding CDC guidelines. We develop and estimate an equilibrium model of patients' physician choice and physicians' prescribing. Counterfactuals show that patient resorting attenuates policies targeting top prescribers. We propose expert-targeted policies that eliminate sorting on preferences exceeding guidelines while preserving care for patients with clinical needs.

Keywords: sorting, opioid prescription, expert-targeted policy

*We thank Abby Alpert, Michael R. Baye, Andrew Butters, Shengmao Cao, Yiqun Chen, Luca Flabbi, Matt Grennan, Paul Grieco, Anran Li, Ethan Lieber, Haizhen Lin, Mark Roberts, Bradley Setzler, Ashley Swanson, Molly Schnell, Vini Singh, Hitoshi Shigeoka, Amanda Starc, Mo Xiao, participants of the 2026 IIOC, 2026 SUFE Empirical IO Conference, and the IU Kelley Brownbag for helpful comments. We thank Richard Meraz for his support in accessing the data used in this paper. This research was supported in part by Lilly Endowment, Inc., through its support for the Indiana University Pervasive Technology Institute. Ng gratefully acknowledges the research support from the National Science and Technology Council of Taiwan.

†Institute of Economics, Academia Sinica. Email: hlng@econ.sinica.edu.tw.
‡Kelley School of Business, Indiana University. Email: ssagl@iu.edu.
§Kelley School of Business, Indiana University. Email: xiangjia@iu.edu.

1 Introduction

In many markets, consumers rely on experts for access to goods and services. Investors turn to asset managers for complex investment opportunities. Patients consult physicians for medical advice and treatment. Choosing the right expert matters. The wrong asset manager may expose investors to financial risk far from their preferences. The wrong physician can significantly affect patients' long-term health. Given these stakes, consumers have strong incentives to sort into experts whose preferences, style, or philosophy align with their own. Market outcomes, therefore, reflect both expert preferences and consumer preferences, linked through consumer sorting. Disentangling these forces is central to policy design. The effectiveness of expert-targeted policies depends on the relative importance of consumer sorting versus expert preferences. Nevertheless, empirical evidence quantifying the contribution of each channel remains extremely limited.

In this paper, we study the importance of consumer sorting versus expert preferences for market outcomes and the implications for expert-targeted policies. We focus on opioid prescribing for chronic pain patients in the United States. Chronic pain affects 60 million adults in the U.S. (Lucas and Sohi, 2024). One-fifth rely on prescription opioids through their chosen physicians (Dahlhamer et al., 2021). The setting is important for two reasons. First, unlike acute pain, chronic pain requires continuous pain management. This exposes patients to a substantially higher risk of opioid misuse. Patient sorting into high-prescribing physicians may significantly exacerbate these risks. Second, despite the prevalence of chronic pain, the role of patient sorting versus physician preference in opioid prescribing to chronic pain patients remains widely understudied. Prior work focuses almost exclusively on quasi-random assignment of patients to physicians, such as acute pain patients admitted through the Emergency Department. Consequently, the literature identified only the role of physicians in prescribing opioids.

Using employer-sponsored health insurance claim data, we show that patient sorting is more than twice as important as physicians' inherent prescribing propensity in explaining the observed dispersion in opioid prescribing across physicians. Patient sorting is almost exclusively driven by factors not concordant with the Centers for Disease Control and Prevention (CDC) clinical guidelines for prescribing opioids (Dowell et al., 2016, henceforth CDC2016). This finding suggests that top physician-targeted policies, such as the Drug Enforcement Administration (DEA)'s prescription license revocation policy, can be improved by accounting for patient resorting. To illustrate the underlying mechanisms, we develop and estimate an equilibrium model of patient sorting and physician prescribing. We use our estimates to simulate several counterfactuals. First, mimicking top expert-targeted policies, we exclude

the top 10% of physicians by risk-adjusted prescription rate. Despite being extreme, this policy reduces the aggregate prescriptions by only 7%. Resorting severely attenuates the exclusion policy. We show that this result is driven by affected patients resorting to non-excluded physicians with similar prescribing propensities. Second, we eliminate resorting. Affected patients may only be randomly assigned a physician or forgo care in response to the same exclusion policy. Without resorting, aggregate prescribing falls by an additional 5 percentage points. However, randomization breaks sorting indiscriminately regardless of patients' medical needs. Third, we instead force physicians to disregard any patient preference exceeding (i) CDC guidelines (CDC2016); or (ii) appropriate medical consensus, such as compassionate prescribing. This removes the incentive to sort on such preferences while preserving treatment for patients with legitimate medical needs. Breaking only this "bad" sorting, aggregate prescriptions drop by 7%–9%, depending on how much medical consensus is deemed appropriate. Lastly, we show that more accurate information on physicians' prescribing propensities facilitates sorting, further attenuating the effect of exclusion policies with resorting.

We establish these findings using health insurance claim data from the Merative MarketScan Research Databases for the New York City Metropolitan Statistical Area over the period 2016–2022. These visit-level data allow us to link patients, their chosen physicians, and any resulting opioid prescriptions. We begin by documenting that physicians' prescription rates are highly dispersed, even after controlling for a rich set of patient severity measures. The top 10% of physicians in terms of prescribing intensity account for approximately 30% of opioid prescriptions. This pattern aligns with the well-documented heterogeneity in physicians' prescribing intensity reported in the literature (e.g., Schnell and Currie, 2018).

Next, we assess the importance of patient sorting for the dispersion of physicians' prescribing decisions. Following the well-known two-way fixed effects approach developed by Abowd et al. (1999), we decompose the variance of physicians' prescription decisions into components attributable to patient fixed effects, physician fixed effects, and patient sorting. There are two identification challenges. First, physician fixed effects are identified from patients who switch physicians. When switching is limited, fixed effects estimates can be severely biased, a problem known as limited mobility bias (Andrews et al., 2012). We follow recent advances in the labor literature to address limited mobility bias by grouping physicians and patients using k-means clustering and estimating group-level fixed effects (Bonhomme et al., 2019; Mourot, 2025). Second, fixed effects estimates may be biased if patients' switching decisions are due to time-varying characteristics unobserved by the econometrician. We provide evidence to mitigate this concern following Card et al. (2018).

Using the two-way fixed effects estimator, we find that patient sorting is three times

as important as physicians' inherent prescribing propensity in explaining the observed dispersion in physicians' opioid prescribing rates. Patient sorting is almost exclusively driven by patient preferences not concordant with the CDC guideline for prescribing opioids for chronic pain (CDC2016). Our findings are consistent with anecdotal evidence from online forums, where patients with strong opioids preferences actively seek and share information about physicians' prescribing propensities, and choose their physician accordingly.[1] Moreover, our findings suggest that policies that target only a subset of physicians may be severely attenuated when ignoring the effect of patient resorting.

To understand how patient sorting interacts with top physician-targeted policies, we develop an equilibrium model of patient choice of physicians and physicians' opioid prescription decisions. The patient choice stage is a static discrete choice demand model in the spirit of Berry et al. (1995). We model patient types as time-persistent tuples that summarize both opioid preferences and information. Patients' information types determine how accurately patients perceive physicians' prescribing propensities. When choosing a physician, patients consider the perceived prescribing probability based on their own opioid preferences and their perceptions about each physician's prescribing propensity. In the prescribing stage, partially altruistic physicians optimally prescribe opioids based on their own prescribing propensities and patients' opioid preferences.

Estimating our model is challenging because it features three dimensions of unobserved heterogeneity: physicians' prescribing propensity, patients' opioid preferences, and patients' information types. To reduce the estimation burden, we estimate our model in three steps using a limited-information maximum-likelihood approach. First, we use our panel on realized opioid prescriptions to estimate both physicians' prescribing propensities and patients' opioid preferences with a two-way fixed effects logit model. This estimation closely mirrors our earlier decomposition of prescription rate variation and follows the grouped fixed effects approach (Bonhomme et al., 2019; Mourot, 2025). This approach addresses both limited mobility and incidental parameter biases. Second, conditional on first-stage estimates, we estimate the patient choice stage by constrained maximum likelihood, imposing a share constraint commonly used in discrete choice demand models (Berry et al., 1995). This step recovers patients' mean utilities for each physician, parameters governing patient preferences for office visits, and the variance of time-invariant patient information types by leveraging the panel structure of the data. Finally, we decompose the estimates for patient opioid preferences into components that are concordant with the CDC guideline, components that

[1] For example, https://www.reddit.com/r/hudsonvalley/comments/1ccgimp/pain_management_doctor_that_will_actually_manage/ (accessed 05/07/2026), and https://www.reddit.com/r/grassvalley/comments/1kn7h0i/doctor_that_prescribes_opiates_in_town_or_nc/ (accessed 05/07/2026)

warrant prescribing based on appropriate medical consensus, and components that are neither in line with the CDC guideline nor with medical consensus. To do so, we use a rich vector of patient-level controls, including patient severity measures and patient diagnosis histories.

We use our estimates to illustrate two points. First, we highlight the mechanism of how patient sorting interacts with top physician-targeted policies, such as the DEA's prescription license revocation policy. Second, we show how they can be improved by accounting for patients' equilibrium resorting. To do so, we consider an extreme version of the top physician-targeted policies by excluding the top 10% of prescribers. We find that even though our policy is extreme, aggregate opioid prescriptions drop by only 7%. This muted response arises because patients do not choose physicians randomly. Instead, our results imply that patients sort toward physicians whose prescribing propensity aligns with their own preferences for opioids. When top physician-targeted policies eliminate these first-best matches, patients optimally resort to the remaining physicians.

Next, we quantify the extent to which resorting attenuates our exclusion policy. In particular, we let patients who visited an excluded physician in the baseline choose only whether to seek care. Conditional on seeking care, patients are randomly assigned to a physician. Under this counterfactual, eliminating the top 10% of physicians by prescription rate reduces aggregate opioid prescriptions by 12%, an additional 5 percentage points drop compared to our previous counterfactual with resorting. These results demonstrate the importance of accounting for demand-side equilibrium responses in supply-side-focused policies targeting a subset of experts.

We further illustrate how top physician-targeted policies can be improved by accounting for patient equilibrium resorting. First, we note that restricting patient choice of physician has important drawbacks. The literature shows that the ability to choose a physician is crucial for patients to remain in care (Hsu et al., 2003). Consistent with these findings, our results imply that when patients are randomized to physicians, most affected patients would prefer to forgo care. Although randomization aggressively reduces aggregate opioid prescriptions, it achieves so by breaking optimal patient-physician matches regardless of whether sorting is based on patient preferences concordant with the CDC guideline or appropriate medical consensus ("good sorting"), or not ("bad sorting"). Second, we show how top physician-targeted policies can leverage "good" sorting while breaking "bad" sorting, preserving patient choice and keeping patients in care. We consider a policy that requires physicians to prescribe only based on patient preferences concordant with CDC guidelines and appropriate medical consensus, but not others.[2] Medical consensus is only partially

[2]Our counterfactual policy is in line with the DEA's objective of reducing medically unjustified opioid

observable, and it is an open question how much prescribing based on medical consensus is appropriate. In particular, there is growing evidence that opioid prescribing for pain as part of the medical consensus may have contributed to drug deaths (Evans and Lieber, 2026). To address this, we bound the policy's effect with two counterfactuals. First, we assume that none of the unobserved patient preferences supports prescribing based on appropriate medical consensus. This shuts down all patient sorting based on unobserved preferences. By eliminating "bad" sorting, this policy reduces aggregate prescribing rates by 9%, a result close to randomization. This constitutes an upper bound on the effectiveness of our policy. Second, we repeat this exercise, but now assume that all unobserved patient preferences support opioid prescribing based on appropriate medical consensus. In this case, aggregate prescriptions would drop by 7%. This is a lower bound on the effectiveness of our policy. Moreover, regardless of our assumptions on the unexplained part of patients' opioid preferences, we show that the number of patient visits under this alternative policy closely matches the number of visits when we allowed for both "good" and "bad" sorting. Additionally, we consider the breaking "bad" sorting policy for all patients without excluding any physician from the market. We find that this policy reduces aggregate prescriptions by up to 29 percentage points, depending on the assumption on appropriate medical consensus.

Our analysis contributes to a growing literature on the drivers of prescription opioid consumption and misuse in the U.S. Prior research has broadly focused on either the supply side or the demand side drivers. On the supply side, research has identified aggressive marketing of OxyContin as a key driver of the opioid crisis (Alpert et al., 2021), highlighted the role of retail pharmacies in opioid distribution and diversion (Janssen and Zhang, 2023; Gui et al., 2024), and studied the importance of physician characteristics, including training, practice style, beliefs about drug diversion, and market competition (Schnell and Currie, 2018; Staiger et al., 2022; Eichmeyer and Zhang, 2022; Currie et al., 2023; Schnell, 2025; Evans and Lieber, 2026). Place-based factors, such as the presence of "pill mill" pain clinics and regulatory policies, have also been shown to matter for risky opioid use (Finkelstein et al., 2025). On the demand side, individuals' economic circumstances have been linked to risky opioid use (Carpenter et al., 2017; Hollingsworth et al., 2017; Case and Deaton, 2021). We contribute to this literature by connecting the supply and demand sides by highlighting the role of patient sorting in shaping opioid prescribing. We show that patient sorting exacerbates observed variation in physician prescribing intensity and helps explain opioid misuse.

prescriptions. To do so, the DEA aims to identify physicians who do not prescribe based on patients' medical needs and revoke their license. However, it is rarely enforced in practice (Office of the Inspector General, 2019).

Next, this paper contributes to research on supply side policies that aim to limit access to opioids by studying the demand side response to such policies within the primary prescription market. A large body of research has studied the effects of various supply side policies such as Prescription Drug Monitoring Programs (Sacarny et al., 2016; Alpert et al., 2024; Sacarny et al., 2025). Some studies have shown that certain supply side restrictions invite demand side substitution to illicit drugs (Alpert et al., 2018; Evans et al., 2019; Soliman, 2025). We instead highlight a different demand-side response: patient resorting once physicians face restrictions, occurring within the primary market. Resorting in the primary market is important because exposure to prescription opioids facilitates the transition to illicit opioid use (Eichmeyer and Zhang, 2022), with 80 percent of heroin users reporting prior prescription opioid misuse (Schnell, 2025). To study this mechanism, we develop an equilibrium model that jointly captures patient choice and physician prescribing. The model allows us to illustrate how patient sorting may affect the effectiveness of physician targeted policies and how to leverage sorting to improve such policies. Modeling sorting is novel. Prior work has relied on settings with quasi-random assignment of patients to physicians, such as acute pain patients admitted through the Emergency Department or physician exit, to plausibly rule out selection caused by patient sorting (Barnett et al., 2017; Eichmeyer and Zhang, 2022; Staiger et al., 2022).

Finally, we add to an emerging literature in health economics relying on two-way fixed effects estimators from the labor literature (for example, Abowd et al. (1999); Bonhomme et al. (2019)) to understand observed differences in care utilization or outcomes (Finkelstein et al., 2016; Molitor, 2018; Mourot, 2025). Finkelstein et al. (2016) use a "mover design" approach combined with a two-way fixed effects estimator to decompose the variance of health utilization into a patient component and a geographic area component. We contribute to the literature in two ways. First, we focus on more granular determinants of opioid prescribing behavior. In particular, we decompose the observed variance of physicians' prescribing behavior into a patient component, a physician component, and patient sorting. Second, we show how patient sorting can exacerbate observed differences in physician prescription rates. We find evidence for strong positive sorting. This implies that patients with stronger opioid preferences disproportionately seek care from higher prescribing physicians. Consequently, the observed dispersion in prescription rates exceeds the underlying differences in inherent physician prescribing intensities.

The rest of the paper is structured as follows. Section 2 describes the setting and the data. Section 3 provides evidence of patient sorting. Section 4 presents a model of patient sorting and physician prescribing. Section 5 discusses model identification. Section 6 presents the estimation procedures and results. Section 7 illustrates how patient sorting interacts with

top physician-targeted policies through counterfactual exercise. Section 8 concludes.

2 Setting and Data

Our setting is opioid prescribing for chronic pain in outpatient office visits. We first present background information regarding the setting and then discuss the data for the analysis.

2.1 Setting

We begin by discussing chronic pain in the United States. According to the CDC, chronic pain is defined as an episode of pain that lasts at least 3 months (CDC2016). In 2023, around 24% of the adult population of the United States suffered from chronic pain (Lucas and Sohi, 2024). Chronic pain is not a condition of old age. It affects adults of all age groups (Dahlhamer et al., 2021). Due to its nature, chronic pain often leaves affected individuals incapacitated and unable to function properly without treatment. For around 22% of adults in the U.S. suffering from chronic pain, treatment involves prescription opioids for pain management (Dahlhamer et al., 2021). The use of prescription opioids for pain management is the most common among working-age adults (Dahlhamer et al., 2021). According to the National Center for Health Statistics, adults aged 45-64 years exhibit the highest rates of prescription opioid use (Dahlhamer et al., 2021). Because most of the affected patients are of working age, almost 2 in 3 payments for opioid prescription collections are covered by private insurance (Schnell, 2025).

In the United States, patients must obtain prescription opioids by prescription from a licensed healthcare provider. About 4 in 5 prescription opioids in the United States are prescribed by physicians (Schnell, 2025). In outpatient settings, patients can generally choose their own physician and can switch at any time, with their choice set limited by their insurance network. Once a patient has chosen a physician, the physician examines the patient and decides whether to prescribe opioids or not. If prescribing opioids, the physician also decides on the daily dose and the number of days of supply. After that, patients have to see a physician to obtain a new prescription.

Prescription opioids are a common part of pain management for both acute and chronic pain in the United States. The use of prescription opioids for *acute pain* for short periods of time is widely accepted as medically justified (CDC2016). However, the use of prescription opioids for *chronic pain* management poses substantial health risks, while there is no clear evidence of any health benefits for patients (Volkow and McLellan, 2016). Metastudies have found an association of adverse health outcomes, such as increased risk of opioid abuse and

dependence, all-cause mortality, and myocardial infarction, with long-term opioid treatment for chronic pain (Chou et al., 2020). There is also growing evidence that long-term prescription opioids may make it harder to effectively manage pain in chronic pain patients. Long-term opioid use can result in opioid induced hyperalgesia in patients, a condition of heightened pain sensitivity that may extend beyond the areas initially affected by chronic pain for which opioids were prescribed (Lee et al., 2011). As a result, apart from acute pain, the CDC Guidelines for Prescribing Opioids mention only three coexisting conditions that may warrant prescribing opioids to chronic pain patients: cancer-related pain treatment, sickle cell disease, and palliative or end-of-life care (CDC2016; CDC2022).[3] Outside of these conditions, prescription opioids should not be used as the first line of treatment, but the decision on whether to use prescription opioids for pain management is generally left to the physician and patient (CDC2016). This ambiguity leaves the eventual decision to use prescription opioids as a course of treatment to the discretion of physicians and patient preference.

2.2 MarketScan Data

The majority of patients who receive prescription opioid treatment for chronic pain are privately insured working-age adults. We therefore use the Merative MarketScan Commercial Database, which provides individual-level, de-identified healthcare claims for millions of enrollees and their dependents in employer-sponsored health insurance plans in the United States.

We construct two data sets for the New York City Metropolitan Statistical Area. We restrict the sample to patients with a primary diagnosis of chronic pain recorded in office visits who received at least one opioid prescription between 2016 and 2022.[4] The first data set consists of all office visits of sample patients linked to their prescription collections. We refer to this data set as the *visit data set*. The second data set augments the visit data with

[3]The guidelines place the three conditions outside its scope, declining to advise against opioid treatment for them. For acute pain, the guidelines recommend the lowest effective dose for no longer than needed. This narrow scope marks a sharp break with earlier practice. In the pre-2016 era, following the medical profession's recognition of pain as the "fifth vital sign" in the 1990s and the Joint Commission's 2001 standards requiring pain assessment in all patients, high patient-reported pain scores were a strong trigger for opioid prescription even for chronic pain patients (Mularski et al., 2006; Baker, 2017). CDC2016 reset the guidance, recommending against opioids as the first-line treatment for chronic pain. CDC2022 retained this stance.

[4]We use the Chronic Conditions Warehouse (CCW) algorithm for "Chronic Pain," revised February 2026. Restricting to patients who received at least one opioid prescription ensures our sample captures relevant patients for whom opioids were a treatment option at some point during their chronic pain episode. Section 3 shows that the correlation between patient preferences and the prescribing propensity of their chosen physician is unchanged when we include patients who never received an opioid prescription during the sample period.

time periods in which a patient chose not to see any physician. We refer to this data set as the *choice data set*. We use the visit data set to establish patient sorting for our motivating evidence and in our structural model estimation, and the choice data set solely for structural model estimation.

Step I: Constructing the visit data set. We link prescription collections to office visits by combining the Outpatient Services Table and the Outpatient Pharmaceutical Claims Table from the MarketScan Commercial Database. Following the literature, we assign a prescription opioid collection to an office visit if (1) the collection occurred within seven days after the office visit, and (2) it was the patient's first opioid collection within that window rather than a refill (Ding and Liu, 2021; Alpert et al., 2025). The Outpatient Services Table records the National Provider Identifier (NPI) of the servicing physician, so this step also links each prescription collection to the prescribing physician. NPI is recorded for every visit in our sample.

Next, we determine whether each linked prescription is for an opioid using the Merative Micromedex RED BOOK, which provides drug information at the National Drug Code (NDC) level. Following the literature, we classify a prescription as an opioid if the drug is listed as an opioid analgesic under the American Hospital Formulary Service (AHFS) Therapeutic Class scheme and as a Schedule II or III substance by the DEA (Staiger et al., 2022; Alpert et al., 2024).[5] This step yields a binary indicator for whether an opioid was prescribed at a given office visit.

We construct patient-level characteristics to capture patients' suitability and preferences for opioid prescriptions. For each visit, we observe patient age, sex, up to four ICD-10 diagnosis codes, and spending in both outpatient and inpatient settings. First, we follow CDC guidelines to identify medical conditions that may qualify for prescription opioids in the treatment of chronic pain, and label these guideline-concordant conditions (CDC2016; CDC2022). They are cancer-related pain, sickle cell disease, and palliative or end-of-life care.[6] Second, we identify diagnoses and risk factors for which the CDC guidelines recommend against opioids or urge caution, and label these guideline-non-concordant conditions. These conditions are sleep-disordered breathing (sleep apnea), renal or hepatic insufficiency, mental health conditions, opioid misuse, substance use disorders, non-fatal overdose, and employment in construction or transportation industries (CDC2016; CDC2022).[7] Appendix

[5]The AHFS opioid analgesics category code is 28:08.08.

[6]In our sample, only 91 of 111,671 visits (0.08%) have acute pain as a secondary diagnosis alongside a primary chronic pain diagnosis. We therefore do not include acute pain as a patient characteristic, and we have verified that adding it as a time-varying characteristic leaves all reported estimates unchanged.

[7]Mental health conditions include anxiety, posttraumatic stress disorder, and depression. Occupations

A provides the corresponding diagnosis codes. We code all of these conditions as patient-specific. Each condition is a time-invariant indicator for a patient which equals one if the diagnosis is recorded at any of a patient's chronic pain office visits, or among their outpatient or inpatient visits in the 12 months preceding any such visit. We do so because these conditions are persistent, so their absence from any single claim does not necessarily imply the patient does not have them. Third, we construct measures of patient health status from each patient's outpatient and inpatient medical history in the 12 months preceding the office visit. They include counts of prior inpatient, outpatient, and emergency department visits; total prior out-of-pocket spending; and the Charlson Comorbidity Index (Charlson et al., 1987).[8] Lastly, we construct per-day out-of-pocket (OOP) costs for prescription opioids. Because OOP costs are unobserved when no opioid is prescribed, we use the expected OOP cost per day, defined as the mean OOP cost within a plan-year-month group among observed users. We provide more details in Appendix C.

Lastly, we restrict the sample to patients with at least two visits and physicians with at least two patient visits during the sample period, to estimate patient and physician fixed effects. Table 1 summarizes the visit data set.

Step II: Constructing the choice data set. For our choice data set, which we use only to estimate the structural model, we supplement the above visit data with months when a patient could have seen a physician but did not.[9] For each month, we define the population of active patients (the set of potential patients) as those with at least one office visit in the corresponding year-quarter. When we observe a patient visiting a physician in a given month, we record their choice.[10] When we do not observe an active patient visiting any physician in a month, we record the outside option of not seeing a physician. For those months, we set the patient's characteristics equal to those of the closest month in the same year-quarter.

warranting caution include driving, using heavy equipment, climbing ladders, working at heights or around moving machinery, or working with high-voltage equipment (CDC2016).

[8]We use the 17 categories of the Charlson/Deyo variant (Deyo, 1992), with the more recent weights from Schneeweiss et al. (2003). See Appendix B for details.

[9]We use a one-month period because 79% of patient-months involve one visit. The modal days of supply per prescription is 30 days, and only 0.2% of prescriptions exceed 30 days.

[10]21% of patient-months involve more than one visit. For these cases, we collapse multiple visits into a single patient-month observation if they are all with the same physician (about 80% of these cases), and we record an opioid prescription if any of the visits resulted in one. If a patient has multiple visits with different physicians, we retain the visit with the physician who has the highest number of prescriptions in that month across all patients. If there is a tie in the total number of prescriptions, we retain the visit with the physician who has the highest visit volume in that month across all patients. We drop the remaining patient-months in which a patient saw more than one physician during that month.

Table 1: Summary Statistics of Chronic Pain Patients

Variable	Mean	SD	p25	p50	p75
Panel A: Guideline-Concordant Conditions					
Cancer-Related Pain	0.02	0.14	0.00	0.00	0.00
Sickle Cell Disease	0.01	0.08	0.00	0.00	0.00
Palliative or End-of-Life Care	0.01	0.09	0.00	0.00	0.00
Panel B: Guideline-Non-Concordant Conditions					
Sleep-Disordered Breathing	0.12	0.33	0.00	0.00	0.00
Renal or Hepatic Deficiency	0.03	0.18	0.00	0.00	0.00
Mental Health Conditions	0.37	0.48	0.00	0.00	1.00
Opioid Misuse	0.07	0.25	0.00	0.00	0.00
Substance Use Disorder	0.16	0.36	0.00	0.00	0.00
Nonfatal Overdose	0.04	0.20	0.00	0.00	0.00
Work in Construction or Transportation	0.06	0.25	0.00	0.00	0.00
Panel C: Patient Severity and Constraints					
Age (years)	48.63	11.30	40.65	51.00	58.00
Female	0.53	0.50	0.00	1.00	1.00
Charlson Comorbidity Index	1.62	2.00	0.00	1.00	2.20
Counts of Prior Inpatient Visits	0.15	0.45	0.00	0.00	0.00
Counts of Prior Outpatient Visits	64.05	69.25	22.25	44.82	81.64
Counts of Prior ED Visits	0.50	1.33	0.00	0.00	0.57
Prior OOP Spending ($)	1462.39	1772.97	346.48	885.75	1972.49
OOP Cost Per Day ($)	1.01	1.44	0.71	0.83	1.17
Observations					
Number of Patients			15,019		
Number of Office Visits			111,671		

Notes: This table summarizes the visit data set. Guideline-concordant conditions are those that may qualify a patient for prescription opioids in the treatment of chronic pain under the CDC guidelines. Guideline-non-concordant conditions are those for which the guidelines recommend against opioids or urge caution (CDC2016; CDC2022). Appendix A lists the corresponding diagnosis codes. All variables are reported at the patient level. Each condition in Panels A and B equals one if the diagnosis is recorded at any of a patient's chronic pain office visits, or among their outpatient or inpatient visits in the 12 months preceding any such visit. Characteristics in Panel C are averages across a patient's chronic pain office visits. We restrict the sample to patients with at least two visits and physicians with at least two patient visits during the sample period.

For our second data set, we also construct physician-level variables. Based on patients' choices among physicians and the outside option, we compute each physician's market share in each market (i.e., year-month). Market share is defined as the number of patients observed choosing that physician that month, divided by the total number of active patients in that month. A common problem in defining market shares based on observable choices is that some options (physicians) may not be chosen every month. When this happens, we follow the standard approach in the discrete choice literature and assume that this physician is inactive that month. We interpret this as the possibility that physicians might close their offices periodically for vacation. Table 2 summarizes the data set used for structural estimation.

Table 2: Physician Market Share Summary Statistics

Variable	Mean	SD	p25	p50	p75
Physician Market Share (%)	0.07	0.11	0.02	0.03	0.06
Inside Share in a Market (%)	50.88	2.86	49.37	51.07	52.71
Observations					
Number of Markets			84		
Number of Physicians			10,355		
Number of Physician-Market			65,414		

Notes: A market is an NYC MSA year-month. The 2016–2022 sample yields 84 markets. A physician's market share is the number of patient visits to that physician in a market divided by market size, expressed in percentage points. The unit of observation in the first row is a physician-market. The inside share is the sum of physician market shares within a market. The unit of observation in the second row is a market.

3 Establishing Patient Sorting

In this section, we first document the heterogeneity in opioid prescription rates among physicians. We then decompose the variance of the prescribing decision into patient opioid preference, physician prescribing propensity, and patient sorting. Finally, we extend the analysis beyond prescribing to opioid misuse, examining the extent to which patient sorting may contribute to risky prescription opioid use.

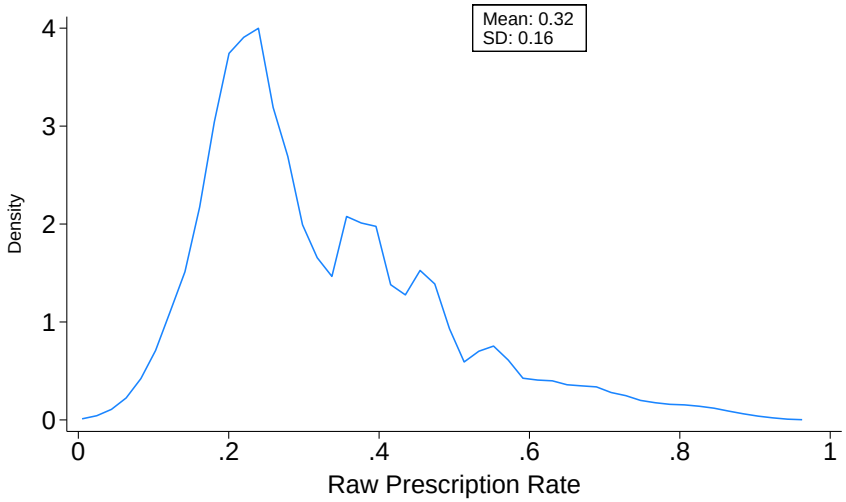
Cross-Physician Variation in Prescription Rates. Our primary outcome is an indicator equal to one if an office visit results in an opioid prescription, and zero otherwise. We estimate physician-level raw prescription rates by regressing this indicator on physician fixed effects,

$$\text{prescribe}_{it} = a_{j(i,t)} + \varepsilon_{it}, \tag{1}$$

where prescribe_{it} equals one if patient i receives an opioid prescription at visit t, $a_{j(i,t)}$ denotes the fixed effect for physician j, capturing physician j's prescription rate, and ε_{it} is an error term. As is well known, fixed effects may be noisily estimated for physicians with few patient visits. Therefore, we follow Chandra et al. (2016) and apply an empirical Bayes shrinkage estimator that shrinks imprecisely estimated fixed effects towards the overall mean.

Figure 1 plots the density of the empirical Bayes–adjusted physician fixed effects $\hat{a}_{j(i,t)}$. Raw prescription rates vary substantially across physicians. The standard deviation is one-half of the mean prescription rate. This finding is consistent with prior evidence documenting large differences in drug prescribing intensity across providers (Schnell and Currie, 2018; Currie and MacLeod, 2020; Currie et al., 2023; Badinski et al., 2023). The distribution is also right-skewed with a long upper tail, indicating that a small share of physicians prescribe opioids at rates far above the average.

Figure 1: Physician Prescription Rate



Notes: This figure plots the distribution of the empirical Bayes–adjusted physician fixed effects from Equation (1).

Next, we investigate the determinants of the observed dispersion in opioid prescription rates. Prescription rates may vary due to differences in the severity of the patient pool. To assess whether differences in opioid prescription rates mainly reflect differences in patient severity, we risk-adjust physicians' prescribing decisions by controlling for a rich set of patient observables. Our data features three categories of patient observables: (i) patient health status and patient-specific constraints, including patient age, sex, the Charlson comorbidity index, counts of historical inpatient, outpatient, and emergency-department vis-

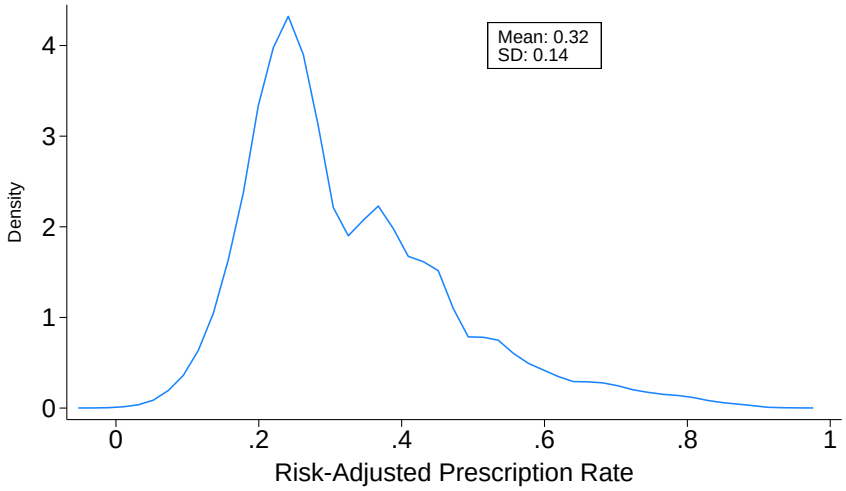
its, historical out-of-pocket spending, and OOP per day; (ii) patient medical conditions for which opioid use may be clinically warranted under the CDC's opioid prescribing guidelines (CDC2016; CDC2022), which in our data are cancer-related pain, sickle cell disease, and palliative or end-of-life care; and (iii) medical conditions and risk factors for which the CDC recommends caution for prescribing opioids for chronic pain, including diagnoses of sleep-disordered breathing, renal or hepatic insufficiency, mental health conditions, opioid misuse, substance-use disorder, nonfatal overdose, and an indicator for employment in construction or transportation industries.

We construct risk-adjusted prescription rates for each physician by regressing the prescription indicator on physician fixed effects, $a_{j(i,t)}$, the full set of patient observables described above, denoted by X_{it}, as well as year fixed effects, τ_t, to absorb any potential common time-varying factors such as policy changes and guideline updates,

$$\text{prescribe}_{it} = a_{j(i,t)} + X_{it}b + \tau_t + \epsilon_{it}. \tag{2}$$

We again apply the empirical Bayes shrinkage procedure following Chandra et al. (2016) and plot the density for the estimated physician fixed effects $\hat{a}_{j(i,t)}$ in Figure 2.

Figure 2: Risk-adjusted Prescription Rate



Notes: This figure plots the distribution of the empirical Bayes–adjusted, risk-adjusted physician fixed effects from Equation (2).

Figure 2 shows that, relative to Figure 1, controlling for patient observables and year fixed effects largely preserves the dispersion in prescription rates. In other words, the prescription rate heterogeneity across physicians is not primarily driven by patient observables or common time trends.

Patient Sorting. We decompose the variance of raw prescription decisions into components attributable to patient opioid preference, physician prescribing propensity, and patient sorting into physicians. Our variance decomposition approach uses the well-known two-way fixed effects estimator of Abowd et al. (1999), but is applied to prescription decisions. We further distinguish sorting based on CDC guideline-concordant versus non-concordant preferences.

To decompose the variance of the prescription decision, we consider the following regression equation,

$$\text{prescribe}_{it} = \eta_{j(i,t)} + \theta_i + X_i\beta + \tau_t + \varepsilon_{it}, \tag{3}$$

where $\eta_{j(i,t)}$ represents the prescribing propensity of the chosen physician, θ_i represents unobserved patient opioid preference, X_i is the time-invariant version of patient characteristics used for risk adjustment, τ_t are year dummies, and ε_{it} is an error term.[11] Note that in Equation (3), β and θ_i are not separately identified because θ_i and X_i are both patient-specific. Therefore, we define $\Theta_i \equiv \theta_i + \beta X_i$, treat Θ_i as patient fixed effect for estimation in the spirit of Abowd et al. (1999), and rewrite Equation (3) as

$$\begin{aligned}\text{prescribe}_{it} &= \eta_{j(i,t)} + \underbrace{\theta_i + X_i\beta}_{\equiv\Theta_i} + \tau_t + \varepsilon_{it} \\ &= \eta_{j(i,t)} + \Theta_i + \tau_t + \varepsilon_{it}.\end{aligned} \tag{4}$$

Obtaining unbiased fixed effects in an Abowd et al. (1999) style model is challenging because physician fixed effects are identified only from patients who switch physicians. When switching is limited, fixed effects estimates can be severely biased, a problem known as limited mobility bias (Andrews et al., 2012; Bonhomme et al., 2023). Recent advances in the labor literature show that this bias can be mitigated using a two-step estimator that first groups firms and then estimates group-level fixed effects (Bonhomme et al., 2019). We adapt this estimator to the setting of opioid prescribing. Following Mourot (2025), in the first stage, we group physicians and patients using k-means clustering based on physician-level and patient-level risk-adjusted prescription rates, respectively.[12] Appendix D provides more details on the grouping algorithm. As in Mourot (2025), we use five groups for physicians and five

[11] Most patient observables exhibit little within-patient variation over time. To address this issue, we use patient-level averages and treat them as time-invariant. In Appendix Table E1, we also estimate a version with time-varying X_{it}. The results are highly similar in terms of the relative importance of each component in explaining variation in prescription rates.

[12] Mourot (2025) groups surgeons and hospitals, whereas we group patients and physicians.

groups for patients. We assess robustness by increasing the number of groups. In the second stage, we estimate Equation (4) using group-level physician and patient fixed effects.

After obtaining estimates for group-level fixed effects, $\hat{\eta}_{j(i,t)}$ and $\hat{\Theta}_i$, we further decompose $\hat{\Theta}_i$ into guideline-concordant and guideline-non-concordant preference components. To do so, we regress patient fixed effects estimates, $\hat{\Theta}_i$, on our full vector of non-time-varying patient observables X_i,

$$\hat{\Theta}_i = \underbrace{X_{1i}\beta_1}_{\text{guideline-concordant}} + \underbrace{X_{2i}\beta_2 + \theta_i}_{\text{guideline-non-concordant}}, \tag{5}$$

where X_{1i} contains the guideline-concordant conditions, X_{2i} contains all remaining guideline-non-concordant patient observables, and θ_i is an error term representing unobserved guideline non-concordant preferences.[13]

We next report the results of our decomposition of the variance of the prescribing decision in Table 3. Following standard practice in the labor literature (Kline, 2024), we scale each component of the decomposition by $var(\text{prescribe}_{it}^E)$, the variation in prescribing that can be explained by either the patient or the physician side factors.[14] For the baseline specification in Table 3, we use five groups each for patients and physicians, and restrict the sample to patients with at least two visits and physicians with at least two patient visits during the sample period. We also report results using alternative numbers of groups, as well as a robustness check following Finkelstein et al. (2016) that includes fixed effects for the timing relative to a patient's switch of physician.[15]

Our results in Table 3 demonstrate that in all specifications, there is strong positive sorting between patients and physicians. In our baseline specification, we find that sorting is overwhelmingly due to guideline-non-concordant preferences for opioid prescriptions. Patient sorting based on guideline-concordant preferences explains almost none of the variation in prescription rates, whereas sorting based on guideline-non-concordant preferences accounts for 33%. Perhaps surprisingly, physicians' intrinsic prescribing propensity explains only about 15% of the variation in prescribing decisions. Our results imply that patient sorting accounts for more than twice as much of the explained variation in prescribing as

[13] Examples of factors captured by θ_i are pain level, which has been used for prescribing based on medical consensus, or any opioid preference exceeding medical necessity that physicians may consider when prescribing.

[14] According to Kline (2024), this normalization ensures the comparability of our results with any future related studies.

[15] To include relative timing fixed effects, we restrict the sample to patients who switch physicians at most once, ensuring relative timing is well defined. This restriction retains 30% of observations. Relative timing indicators are normalized to zero for never-switchers.

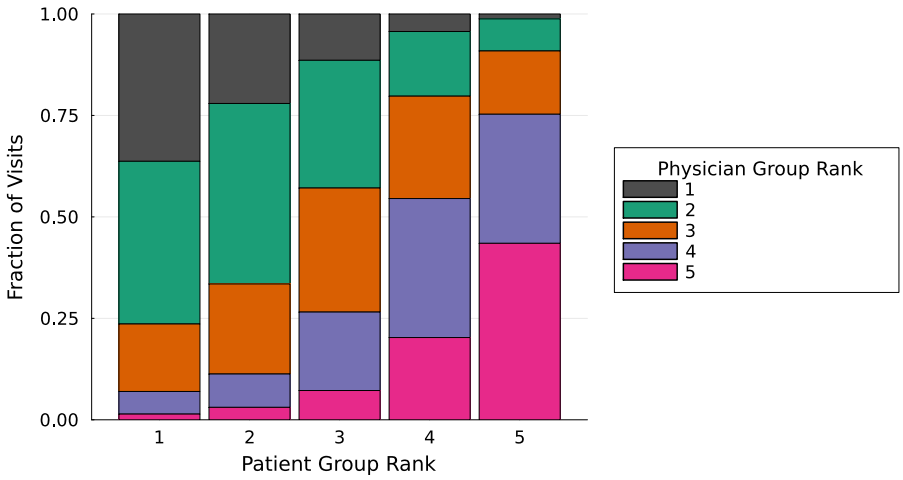
physicians' own prescribing propensities (33% versus 15%). Consistent with this finding, the correlation coefficient between patients' guideline-non-concordant preferences and their chosen physicians' prescribing propensities is as high as 0.59.

Moreover, Appendix Table E1 presents three sets of robustness checks: (i) alternative sample definitions, including one that adds patients who never received an opioid prescription during the sample period; (ii) specifications that include time-varying patient characteristics X_{it}; and (iii) an empirical Bayes–adjusted estimator for the patient and physician fixed effects before they are used as inputs to the k-means grouping. All robustness checks yield qualitatively highly similar results.

We further illustrate positive sorting in Figure 3 following Mourot (2025). The figure shows that patients with the lowest opioid preference (patient group 1) make more than three-quarters of their visits to physicians whose prescribing propensities are the lowest or second lowest (physician groups 1 and 2). Conversely, patients with the highest opioid preference (patient group 5) make more than three-quarters of their visits to physicians in the top two prescribing groups (physician groups 4 and 5), with 44% concentrated in the highest group alone. This shift is evidence of positive sorting. Across the full sample, visits are distributed across the five physician groups as 18%, 31%, 21%, 17%, and 13%, from the lowest to the highest prescribing propensity. Under random matching of patients and physicians, every bar in the figure would reproduce this same composition. Physician groups 1 and 2 would therefore together receive 49% of visits in every patient group. In the data, however, they receive more than three-quarters of visits from group-1 patients. Likewise, physician group 5 would receive 13% of visits in every patient group but in fact accounts for 44% of visits from group-5 patients. Appendix Figure F1 shows the same sorting patterns when patients and physicians are divided into 10 groups instead.

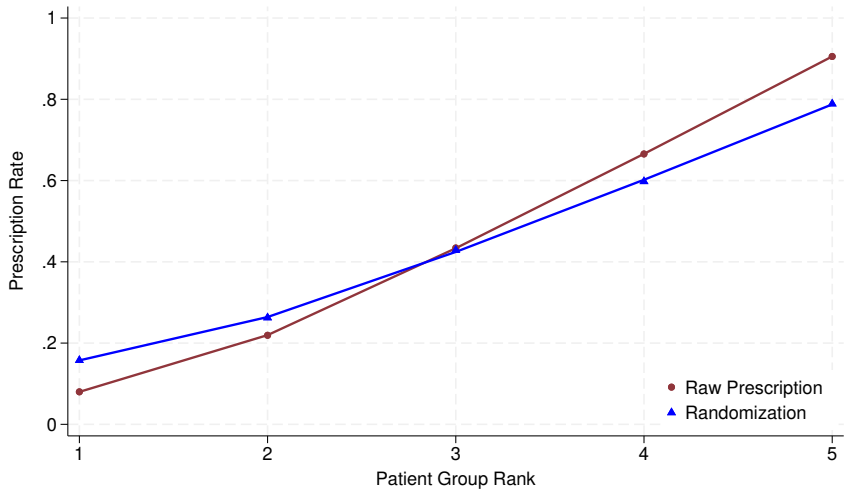
Strong positive sorting implies that the observed prescription rates for each patient group differ from the prescription rates that would be observed if patients chose physicians at random. Figure 4 illustrates this point. Our results show a stark departure of the data from the predicted prescription rates under random assignment. Observed prescription rates for low opioid preference types are considerably lower than the one implied under randomization. This captures the strong underlying incentives for these patients to sort into lower-prescribing physicians. Similarly, observed prescription rates for high opioid preference types are considerably higher than the one implied under randomization. Overall, these results demonstrate that patient sorting amplifies differences in prescription rates across patient groups.

Figure 3: Patient Sorting



Notes: This figure plots, for each of the five patient groups, the distribution of visits across the five physician groups. Patient and physician groups are formed from their respective fixed effects estimates and ranked from group 1 (lowest type) to group 5 (highest type).

Figure 4: Raw Prescription Rates Versus Prescription Rates under Random Assignment



Notes: This figure compares observed prescription rates for each patient group with the prescription rates that would be observed if patients chose physicians at random. When randomly assigning a physician to each visit, we weight physicians by visit share. Patients are partitioned into five groups ranked by their estimated fixed effect Θ_i, shown on the horizontal axis. Red circles plot the average raw prescription rate, and blue triangles plot the average prescription rate under random assignment.

Table 3: Variance Decomposition of Prescription Decision (%)

	Baseline	**Alt. Number of Groups**				**Adding Relative Timing FE**
	$K = 5$ ≥ 2 visits/patient ≥ 2 visits/physician	$K = 10$	$K = 15$	$K = 20$	$K = 50$	Following Finkelstein et al. (2016)
Patient						
$\frac{\text{Var}(X_{2i}\beta_2 + \theta_i)}{\text{Var}(\text{prescribe}^E)}$	52.36	52.20	52.37	52.70	52.78	46.13
$\frac{\text{Var}(X_{1i}\beta_1)}{\text{Var}(\text{prescribe}^E)}$	0.05	0.06	0.07	0.06	0.06	0.16
Physician						
$\frac{\text{Var}(\eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$	14.78	14.90	14.77	14.74	14.73	17.50
Sorting						
$2 \times \frac{\text{Cov}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$	32.98	33.01	32.99	32.68	32.61	36.44
$2 \times \frac{\text{Cov}(X_{1i}\beta_1, \eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$	-0.03	-0.02	-0.03	-0.03	-0.03	-0.01
Correlation						
$\text{corr}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})$	0.59	0.59	0.59	0.59	0.58	0.64
$\text{corr}(X_{1i}\beta_1, \eta_{j(i,t)})$	-0.02	-0.01	-0.02	-0.02	-0.01	-0.00

Notes: This table decomposes the variance of the prescription decision into patient, physician, and sorting components, with each component scaled by $\text{Var}(\text{prescribe}^E_{it}) \equiv \text{Var}(X_{1i}\beta_1 + X_{2i}\beta_2 + \theta_i + \eta_{j(i,t)})$, the variation in prescribing explained by patient- and physician-side factors. The last column reports a robustness check following Finkelstein et al. (2016) that adds fixed effects for the timing relative to a patient's switch of physician. These indicators are normalized to zero for non-switchers. The sample is restricted to patients who switch at most once so that relative timing is well defined, which accounts for 30% of observations.

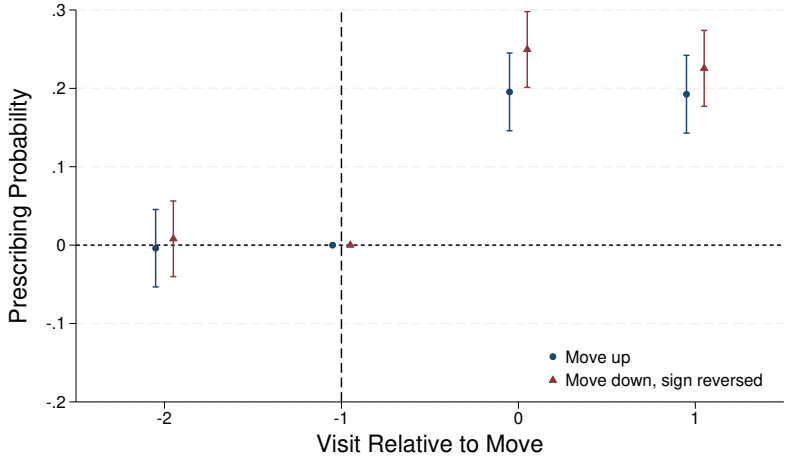
Threats to Identification. Including patient fixed effects, which absorb all time-invariant unobserved patient characteristics, addresses many potential confounders. However, fixed effects estimates may be biased if patient switching is correlated with unobserved time-varying factors contained in ε_{it}. In this case, switchers differ from non-switchers in their prescribing trends around the time of the switch. For example, patients may switch physicians in response to a health shock that directly affects their opioid preferences. To address this concern, we follow Finkelstein et al. (2016) and include visit fixed effects relative to a physician switch, flexibly accounting for potential differences around switching. We report the results in the last column of Table 3. The results of this exercise are qualitatively highly similar to our baseline specification.

To further mitigate concerns about endogenous switching, we follow the labor literature to provide additional suggestive evidence that switching is not driven by patients' omitted time-varying unobservables, such as their changing opioid preferences in our setting. In particular, Card et al. (2018) shows that if worker switching is exogenous, workers exhibit symmetric wage responses with opposite signs when switching to higher- versus lower-type firms. Analogously, if patient switching is exogenous in our context, patients who switch to more liberal prescribers should experience changes in prescription probability that are symmetric, with opposite signs, to those of patients who switch to more conservative prescribers. Following Card et al. (2018), we conduct an event study analysis and plot the magnitude of the change in prescribing probability in event time relative to the switch for two patient groups: (i) patients who move from a lower-type physician to a higher-type physician (move-up); and (ii) patients who move from a higher-type to a lower-type physician (move-down). We use patients who move to a physician in the same physician group as the control group. To conduct our event study, we restrict the sample to switchers with two visits before and one visit after the switch.[16] We present our findings in Figure 5. Two results mitigate concerns that the switching is based on unobserved time-varying factors. First, we find no evidence of pre-trends in prescribing behavior prior to the switch, indicating that the switching is not driven by patients' gradual learning or pre-move health deterioration. Second, the post-switch responses are symmetric. The magnitude of the change for move-up patients is not significantly different from that for move-down patients, while the signs of the changes are opposite.

Patient Sorting Contributes to Risky Opioid Use. We established that patient sorting is an important driver of prescription rate heterogeneity. Next, we examine whether and

[16]Because patients in our data are observed for a limited number of visits, requiring three or more pre-switch visits would significantly reduce the sample size. Our event study uses 16,604 observations, or 15% of the full sample.

Figure 5: Magnitude of Prescription Rate Responses to Patient Switching



Notes: This figure presents event study estimates of the change in the prescribing probability around a physician switch, separately for patients who switch to higher-prescribing physicians ("move up") and to lower-prescribing physicians ("move down"), with 95% confidence intervals. The horizontal axis measures the visit relative to the switch, where time 0 corresponds to the visit at which the switch occurs. To make the changes comparable, we plot the negative of the estimated change for downward movers.

to what extent positive patient sorting can explain adverse health outcomes. We revisit the decomposition exercise in Equation (4), replacing the prescription outcome variable with a patient health outcome. To do so, we follow a common approach in the literature and measure patient health outcomes using an indicator for opioid misuse at a visit. We classify a visit as involving opioid misuse if the patient's daily Morphine Equivalent Dose (MED) is 120 mg or greater (Meara et al., 2016; Staiger et al., 2022; Finkelstein et al., 2025). We focus on this health outcome for three reasons. First, it is well defined at the visit level. Second, it provides a more objective measure of opioid misuse than opioid-abuse–related diagnoses coded by the physician. These diagnoses depend on both the physician's willingness to assess and the patient's willingness to undergo testing. Moreover, the literature has documented a strong correlation between opioid misuse and downstream adverse outcomes, including opioid poisonings and opioid addiction (Finkelstein et al., 2025). We provide details on the construction of our measure of daily MED in Appendix G. Based on this definition, 6% of visits in our sample involve risky opioid use.[17]

In Table 4, we decompose the variation in opioid misuse into components attributable to guideline-concordant patient preferences ($X_{1i}\beta_1$), guideline-non-concordant patient preferences ($X_{2i}\beta_2 + \theta_i$), physician prescribing propensity ($\eta_{j(i,t)}$), and patient sorting on both guideline-concordant ($Cov(X_{1i}\beta_1, \eta_{j(i,t)})$) and guideline-non-concordant preferences ($Cov(X_{2i}$

[17] Finkelstein et al. (2025) used Social Security Disability Insurance (SSDI) data in a 20% random sample of Medicare beneficiaries from 2006 to 2019 and identified 3% of visits involve risky opioid use.

$\beta_2 + \theta_i, \eta_{j(i,t)}))$. All specifications yield qualitatively similar results. Moreover, the patterns closely align with the decomposition of opioid prescribing decisions in Table 3. Our results confirm that patient sorting is an important driver of patient health outcomes. Patient sorting explains 20% of the variation in opioid misuse. Our results further demonstrate that patients sort on preferences that do not support prescribing based on CDC guidelines. The correlation coefficient between patient guideline-non-concordant preference type and physician propensity is 0.53. Appendix Table E2 presents robustness checks using alternative sample definitions, time-varying patient characteristics X_{it}, and Bayes shrinkage. All yield qualitatively similar results. These results indicate that patient sorting not only drives high opioid prescription rates, but is also associated with worse health outcomes.

Discussion. Our findings imply that patients do not randomly choose their physicians. Patient sorting plays an important role in explaining cross-physician variation in both prescribing decisions and opioid misuse. Moreover, patient sorting is primarily based on guideline-non-concordant preferences. Our findings have potentially important policy implications. When top-physician-targeted policies exclude a subset of physicians, they break optimal patient-physician matches. In response, affected patients may resort to the remaining physicians according to their opioid preferences. This suggests that top-physician-targeted policies may be improved by explicitly leveraging patient resorting. However, our motivating evidence cannot offer insights into how patients would optimally resort in response to exclusion policies. To understand how patient sorting interacts with exclusion policies, we therefore develop and estimate a structural model that captures patient sorting through patient optimal choice of physician in the following sections.

Table 4: Variance Decomposition of Opioid Misuse (%)

	Baseline	**Alt. Number of Groups**				**Relative Timing FE**
	$K = 5$ ≥ 2 visits/patient ≥ 2 visits/physician	$K = 10$	$K = 15$	$K = 20$	$K = 50$	Following Finkelstein et al. (2016)
Patient						
$\frac{\text{Var}(X_{2i}\beta_2 + \theta_i)}{\text{Var}(\text{misuse}^E)}$	75.43	76.20	77.10	76.79	77.43	72.27
$\frac{\text{Var}(X_{1i}\beta_1)}{\text{Var}(\text{misuse}^E)}$	0.05	0.05	0.06	0.05	0.05	0.03
Physician						
$\frac{\text{Var}(\eta_{j(i,t)})}{\text{Var}(\text{misuse}^E)}$	4.67	4.37	4.00	4.11	3.95	4.54
Sorting						
$2 \times \frac{\text{Cov}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})}{\text{Var}(\text{misuse}^E)}$	19.87	19.41	18.87	19.08	18.60	23.13
$2 \times \frac{\text{Cov}(X_{1i}\beta_1, \eta_{j(i,t)})}{\text{Var}(\text{misuse}^E)}$	0.00	0.00	0.00	0.00	0.00	0.01
Correlation						
$\text{corr}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})$	0.53	0.53	0.54	0.54	0.53	0.64
$\text{corr}(X_{1i}\beta_1, \eta_{j(i,t)})$	0.00	0.00	0.00	0.00	0.00	0.02

Notes: This table decomposes the variance in opioid misuse into patient, physician, and sorting components, with each component scaled by $\text{Var}(\text{misuse}^E_{it}) \equiv \text{Var}(X_{1i}\beta_1 + X_{2i}\beta_2 + \theta_i + \eta_{j(i,t)})$, the variation in opioid misuse explained by patient- and physician-side factors. The last column reports a robustness check following Finkelstein et al. (2016) that adds fixed effects for the timing relative to a patient's switch of physician. These indicators are normalized to zero for non-switchers. The sample is restricted to patients who switch at most once so that relative timing is well defined, which accounts for 30% of observations.

4 Model

Our results in Section 3 suggest that patients choose physicians according to physicians' prescribing propensities. In this section, we develop a static two-stage model of physician-patient interaction that captures patient sorting. In the first stage, the *choice stage*, in each year-month t, patients choose either to see a physician $j \in \mathcal{J}_t$ in their market (year-month) or to select the outside option of not seeking care.[18] Patients' choice is based on their perceptions about each physician's prescribing behavior. In the second stage, the *prescribing stage*, the chosen physician decides whether to prescribe or not. We illustrate our two-stage model for a single month t and a single patient i in Figure 6.

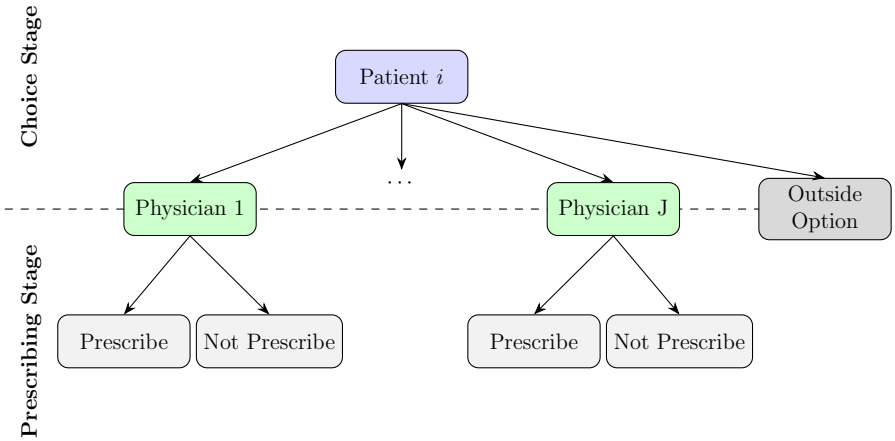


Figure 6: Illustration of two-stage model for each month t

4.1 Indirect Utility and Information Environment

Physician's Indirect Utility Function. In each year-month t, conditional on being chosen by patient i, physician j makes a binary prescribing decision. Physician j's indirect utility from treating patient i is given by

$$V_{ijt} = a_{ijt}\big(\eta_j + \tau_t\big) + U_{ijt}(a_{ijt}) + \varepsilon_{ijt}(a_{ijt}). \tag{6}$$

Here, the binary variable a_{ijt} denotes a physician's prescribing decision. It equals 1 if physician j prescribes an opioid to patient i in month t, and 0 otherwise. η_j is a physician-specific fixed effect capturing physician j's inherent propensity to prescribe opioids. It reflects

[18] We define a market as a year-month because almost 80% of patients have at most one visit in a given month, and the modal number of days of supply per prescription is 30. We include more details on how we handle patients who visit more than once in Section 2.2.

factors such as treatment philosophy, medical school training, and risk tolerance. Physicians with a higher η_j are, ceteris paribus, more likely to prescribe an opioid. We normalize one of the η_j parameters to zero. τ_t denotes yearly fixed effects for physician utility from prescribing an opioid, capturing common shifts in the prescribing environment such as guideline changes and the expansion of prescription drug monitoring programs. $U_{ijt}(a_{ijt})$ represents the patient's utility as a function of the physician's prescribing decision.[19] Lastly, ε is an i.i.d. alternative-specific Type 1 extreme value shock, independent of all other components.

Physicians' Information. Physicians make prescribing decisions by considering both observable patient preferences and their own preferences over treatment. We assume that, through their medical expertise, physicians observe the marginal benefit of prescribing an opioid for patient i, $U_{ijt}(a_{ijt} = 1) - U_{ijt}(a_{ijt} = 0)$. We remain flexible about physician information on the remaining components of patient utility. For example, our model allows for physicians to know only the distribution of the remaining patient preferences or to observe them fully. In addition, we assume that this informational structure is common knowledge to both patients and physicians.

Patient's Indirect Utility Function. Patient i's indirect utility from choosing physician j at time t is

$$U_{ijt}(a_{ijt}) = \delta_{jt} + a_{ijt}\Theta_i + \kappa_i + \lambda_{ijt} + \nu_{ijt}, \tag{7}$$

where δ_{jt} denotes the mean utility patients derive from an office visit to physician j in month t. The mean utility may vary across physicians because they differ in amenities or are subject to physician-specific demand shocks. For example, a physician may have friendlier front-desk staff, offer appointment reminders, or be perceived as having greater expertise in treating chronic pain, regardless of their prescribing behavior. Θ_i measures patient i's inherent preference for opioid prescriptions. We refer to Θ_i as patients' *opioid preference type*. As in Equation (5), we assume that patients' preference types are a linearly separable function of patients' guideline-concordant preferences ($X_{1i}\beta_1$) and guideline-non-concordant preferences ($X_{2i}\beta_2 + \theta_i$). We further assume that the distribution of patients' opioid preference type Θ_i has mean zero, $\mathrm{E}[\Theta_i] = 0$. Given our assumptions, Θ_i represents patient i's marginal utility from an opioid prescription.

Our model also allows patients to differ in the utility they derive from an office visit

[19] In principle, our model could accommodate an altruism parameter for patient utility U_{ijt}. This is equivalent to dividing each component in Equation (6) by the altruism parameter. The physician fixed effect η_j would then capture both altruism and prescribing propensity.

for reasons unrelated to an opioid prescription. First, we allow patients to differ in their inherent utility from an office visit via the patient fixed effect κ_i. Patients with higher κ_i are more likely to seek care in any given month t. Second, patients may prefer different physicians for reasons beyond obtaining an opioid prescription. For example, patients with a high preference for opioids may seek care from more liberal prescribers because of a shared treatment philosophy or because they feel less stigmatized. We capture such factors via a physician-patient match value λ_{ijt}, which has two components,

$$\lambda_{ijt} = \bar{\lambda}\Theta_i(\eta_j + \tau_t) + \epsilon_{ij}. \tag{8}$$

The interaction between patient opioid preference type Θ_i and physician prescribing type η_j, $\bar{\lambda}\Theta_i(\eta_j + \tau_t)$, captures synergies between high/low prescribing physicians and high/low preference patients. The second component, ϵ_{ij}, captures time-persistent, pair-specific match value orthogonal to $\bar{\lambda}\Theta_i(\eta_j + \tau_t)$. For example, a high-preference patient i may nonetheless have a high match value for a low-prescribing physician j because of a referral tie. Conversely, a high-preference patient may have a low match value for a high-prescribing physician because they do not harmonize well. We assume $\epsilon_{ij} \sim N(0, \sigma^2_\lambda)$, such that $\lambda_{ijt} \sim N(\bar{\lambda}\Theta_i(\eta_j + \tau_t), \sigma^2_\lambda)$.

Lastly, ν_{ijt} is an i.i.d. Type 1 extreme value shock. We normalize the value of the outside option of not seeing any physician to $U_{i0t} = \nu_{i0t}$.[20]

Patients' Information. We assume that patients do not have perfect information about physicians' preferences. First, patients cannot observe the alternative-specific preference shock $\varepsilon_{ijt}(a_{ijt})$ because it is realized only during the office visit after a physician was chosen. We assume, however, that patients know the distribution of this shock. Second, patients do not perfectly observe physicians' inherent prescribing propensity η_j when choosing a physician. Instead, we assume that patient i's perceived prescribing type of physician j, denoted by ω_{ij}, equals the physician's true type plus noise,

$$\omega_{ij} = \eta_j + \zeta_{ij}. \tag{9}$$

We refer to the vector $\boldsymbol{\zeta_i}$ as patients' *information type*. A patient's information type captures their uncertainty about a physician's prescribing type η_j. We allow patients to differ in their level of informedness for each physician, and assume that $\zeta_{ij} \overset{\text{i.i.d.}}{\sim} N(0, \sigma^2_\zeta)$. Under this information structure, a patient may overestimate one physician's prescribing propensity while underestimating another's. We interpret this assumption as reflecting id-

[20] We interpret the outside option as including seeking pain relief from alternative sources, such as obtaining opioids from the secondary market.

iosyncratic misperceptions about individual physicians, depending on the amount and source of information a patient has for different physicians. In practice, patients may differ in their ability to search for and interpret informal information sources, such as online forums (e.g., Bluelight, Reddit), word-of-mouth, or prior experiences, which can systematically shape how well informed they are about physicians' willingness to prescribe opioids. By assumption, these misperceptions share a common variance σ_ζ^2. If all patients have an accurate perception of each physician's prescribing type, the variance of ζ_{ij} is zero.[21] Finally, we assume that the information structure is common knowledge to both physicians and patients.

4.2 Equilibrium

This two-stage model is solved by backward induction. We start by describing a physician's optimal prescribing behavior, conditional on being chosen by the patient. We then derive patients' optimal physician choice in the first stage.

4.2.1 Physicians' Optimal Prescribing Choice

Physicians can observe a patient's marginal utility from prescribing, Θ_i. Physicians' indirect utility V_{ijt} as an explicit function of the prescribing decision a_{ijt} is

$$V_{ijt}(a_{ijt}=1) = \eta_j + \tau_t + \delta_{jt} + \Theta_i + \kappa_i + \lambda_{ijt} + \nu_{ijt} + \varepsilon_{ijt}(1), \quad \text{and} \tag{10}$$

$$V_{ijt}(a_{ijt}=0) = \delta_{jt} + \kappa_i + \lambda_{ijt} + \nu_{ijt} + \varepsilon_{ijt}(0). \tag{11}$$

Physician j will choose to prescribe whenever their indirect utility from prescribing an opioid weakly exceeds the indirect utility from not prescribing,

$$V_{ijt}(a_{ijt}=1) - V_{ijt}(a_{ijt}=0) = \eta_j + \Theta_i + \tau_t + \varepsilon_{ijt}(1) - \varepsilon_{ijt}(0) \geq 0. \tag{12}$$

This difference and, therefore, the physician's prescribing decision, is *independent* of the demand-side parameters δ_{jt}, κ_i, λ_{ijt}, the patient's information type $\boldsymbol{\zeta_i}$, and the patient's

[21] Our model abstracts from patient learning of physicians' prescribing propensities. We believe this simplification is reasonable because patient sorting remains stable when we restrict the sample to patient-physician pairs with at least one, two, or three prior shared visits and estimate Equation (4) on each subsample. In Appendix Table E3, the sorting component, $2 \times \frac{\text{Cov}(\Theta_i, \eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$, stays within a tight range across subsamples. If learning were a first-order force within our sample, sorting should be more pronounced among experienced physician-patient pairs, who have had more opportunity to gather information. The absence of such an amplification effect lends support to this model simplification.

i.i.d. Type 1 extreme value shock ν_{ijt}. The physician only considers their own prescribing preferences and patients' marginal utility from prescribing. Physicians do not consider their relative attractiveness to the patient, patient fixed effects from an office visit, their match value, or the patient's information about their prescribing propensities. Consequently, our model is flexible about how much information physicians have about patient preferences other than the marginal benefit from prescribing.

4.2.2 Patients' Optimal Choice of Physician

Patients face a discrete choice among J_t physicians in a market, given their information about each physician. As described in Section 4.1, patients do not observe the realization of physicians' Type 1 extreme value shocks ε. Integrating over the Type 1 extreme value shocks yields the probability that physician j prescribes an opioid to patient i at time t,

$$\rho_{ijt} = \frac{\exp(\eta_j + \Theta_i + \tau_t)}{1 + \exp(\eta_j + \Theta_i + \tau_t)}. \tag{13}$$

However, in addition to only knowing the distribution of ε, patients hold potentially inaccurate perceptions ω_{ij} about each physician's prescribing propensity η_j. Consequently, patients cannot perfectly predict a physician's treatment decision at the time of choosing a physician. Instead, based on their information, patients form a perceived probability of receiving a prescription from physician j,

$$\tilde{\rho}_{ijt} = \frac{\exp(\omega_{ij} + \Theta_i + \tau_t)}{1 + \exp(\omega_{ij} + \Theta_i + \tau_t)}. \tag{14}$$

Given patients' perceived prescribing propensity of physician j, ω_{ij}, the perceived probability of receiving a prescription conditional on a visit, $\tilde{\rho}_{ijt}$, and ν_{ijt} drawn i.i.d. from a Type 1 extreme value distribution, the probability that patient i chooses physician j at time t is

$$s_{ijt} = \frac{\exp\left(\delta_{jt} + \tilde{\rho}_{ijt} \times \Theta_i + \kappa_i + \bar{\lambda}\Theta_i(\omega_{ij} + \tau_t) + \epsilon_{ij}\right)}{1 + \sum_k \exp\left(\delta_{kt} + \tilde{\rho}_{ikt} \times \Theta_i + \kappa_i + \bar{\lambda}\Theta_i(\omega_{ik} + \tau_t) + \epsilon_{ik}\right)}. \tag{15}$$

Lastly, we obtain aggregate equilibrium choice probabilities s_{jt} by integrating over individual choice probabilities s_{ijt}

$$s_{jt} = \int s_{ijt} \mathrm{d}i. \tag{16}$$

5 Identification

The parameter vector of the model is $\{\boldsymbol{\eta}, \boldsymbol{\Theta}, \boldsymbol{\tau}, \boldsymbol{\delta}, \boldsymbol{\kappa}, \bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$. Our identification strategy exploits the fact that the parameter vector can be partitioned into two sets. The first set consists of parameters that govern both the choice stage and the prescribing stage of the model, $\{\boldsymbol{\eta}, \boldsymbol{\Theta}, \boldsymbol{\tau}\}$. The second set consists of parameters that exclusively govern the choice stage, $\{\boldsymbol{\delta}, \boldsymbol{\kappa}, \bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$.

We start with the first set of parameters, $\{\boldsymbol{\eta}, \boldsymbol{\Theta}, \boldsymbol{\tau}\}$, which govern both the prescribing and the choice stage. Importantly, these parameters can be identified solely from variation in the prescribing stage. In particular, identification follows the same logic as the two-way fixed effects model used in our motivating evidence, following Abowd et al. (1999). Physician prescribing propensities $\boldsymbol{\eta}$ are identified from variation in observed prescriptions for the same patient across different physicians. Similarly, the vector of patient opioid preference types $\boldsymbol{\Theta}$ is pinned down from variation in observed prescriptions written by the same physician for different patients. As in Abowd et al. (1999), parameters are identified only for physicians and patients who are connected through patient switching. Note that this requires patients to switch between physicians exogenously. Switching can be based on $\boldsymbol{\eta}$, $\boldsymbol{\Theta}$, and $\boldsymbol{\tau}$ but must be unrelated to physicians' prescribing stage error ε_{ijt}. Since patients cannot observe ε_{ijt} when they make their decision, our model satisfies this condition.[22] The coefficients of the year fixed effects $\boldsymbol{\tau}$ are identified from within-patient variation in observed prescriptions across years, holding fixed both time-invariant patient opioid preference types $\boldsymbol{\Theta}$ and time-invariant physician prescribing propensities $\boldsymbol{\eta}$. As we discussed in Section 2, our data allows us to track the vast majority of patients over multiple years, providing rich identifying variation.

Conditional on the identification of $\{\boldsymbol{\eta}, \boldsymbol{\Theta}, \boldsymbol{\tau}\}$, we now discuss how data from the choice stage of our model identify the remaining parameters $\{\boldsymbol{\delta}, \boldsymbol{\kappa}, \bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$. We start by discussing the identification of $\boldsymbol{\delta}$ and $\boldsymbol{\kappa}$. The physician mean utility parameters $\boldsymbol{\delta}$ and patient fixed effects $\boldsymbol{\kappa}$ are jointly identified from observed choices. $\boldsymbol{\delta}$ are identified from variation in the choice of physician within a month, conditional on seeking care, whereas $\boldsymbol{\kappa}$ are identified from variation across patients in the probability of seeking care from *any* physician. In Section 2, we showed that patients differ widely in their probability of seeking care and their choice of physicians, providing rich identifying variation for both $\boldsymbol{\delta}$ and $\boldsymbol{\kappa}$.

The remaining parameters $\{\bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$ are identified from covariation of patient choice probabilities with physician prescribing propensities. Note that identification of these parameters requires dispersion in physicians' prescribing propensities, which we documented

[22] We provide supporting evidence for our assumption by showing that patients switch physicians exogenously in Section 3.

in Section 3. Patient choice responds to differences in prescribing propensities in two ways. First, through variation in the probability of receiving a prescription, and second, through a higher match value orthogonal to an opioid prescription. Separate identification of $\{\bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$ rests on functional form. In particular, our identification argument exploits how variation in the marginal utility of a prescription affects sorting through these channels.

Incentives to sort based on receiving an opioid prescription are nonlinear in the marginal utility from an opioid prescription $\Theta_i + \tau_t$. Very high or very low values of $\Theta_i + \tau_t$ push prescribing probabilities into the logit tails. In this region, patients will either receive an opioid prescroption from any physician with high probability, or none from any. Similarly, incentives to sort solely to obtain an opioid prescription vanish if the marginal utility of an opioid prescription, $\Theta_i+\tau_t$, is close to 0. Match values, however, enter patients' indirect utility linearly, explaining why patients with very high or very low opioid preference still choose with high or low prescribing physicians. Sorting of patients whose objective prescribing probability is in either tail of the prescribing distribution thus pins down $\bar{\lambda}$. We document evidence for the existence of this channel in our setting in Figure 3. Figure 3 shows that patients with the highest and lowest preference still disproportionately sort into physicians with the highest and lowest prescribing propensity. The dispersion of patient information types σ_ζ is identified from patients whose objective prescribing probability is neither in the tails of the logit, nor have a marginal utility of prescribing close to zero. If these patients systematically and persistently deviate from their objective optimal choice, then $\sigma_\zeta > 0$.

Lastly, σ_λ is identified from patients whose marginal utility from an opioid prescription Θ_i is close to zero. Perceived prescribing propensities ω_{ij} enter patient utility through the perceived probability of receiving a prescription and through the match, both interacted with Θ_i. When Θ_i is nearly zero, patients are indifferent to a physician's prescribing propensity, and their choice is governed solely by the second component of the match value. Dispersion in their persistent physician matches identifies σ_λ.

6 Estimation and Results

In this section, we describe how we estimate our structural model and present the estimation results.

6.1 Estimation

Our estimation procedure closely follows our identification argument. In particular, we partition our model parameters into two sets. We then estimate the choice stage and the

prescribing stage separately using a limited information maximum likelihood estimation procedure.

Estimation proceeds in three steps. In the first step, we estimate physicians' prescribing propensities η_j, patients' opioid preference types Θ_i, and yearly fixed effects $\boldsymbol{\tau}$ using realized prescribing decisions. In the second step, conditional on our estimates for $\{\boldsymbol{\eta}, \boldsymbol{\Theta}, \boldsymbol{\tau}\}$, we estimate physician mean utility δ_{jt}, patient utility from an office visit to any physician κ_i, the standard deviation of the match value σ_λ, and the standard deviation of patient information type σ_ζ, as well as $\bar{\lambda}$, the parameter governing how visit utility varies with the interaction between a patient's opioid preference type and a physician's prescribing propensity. In the final stage, we decompose estimated patients' types Θ_i into guideline-concordant preferences, guideline-non-concordant preferences reflected in observables, and residual unobservable guideline-non-concordant preferences. We elaborate on each estimation step below.

Step I: Estimation of $\{\boldsymbol{\Theta}, \boldsymbol{\eta}, \boldsymbol{\tau}\}$ from the prescribing stage. Note that the prescribing stage depends only on $\{\boldsymbol{\Theta}, \boldsymbol{\eta}, \boldsymbol{\tau}\}$. We can rewrite Equation (13) as the following two-way fixed effects logistic regression model,

$$\log(\frac{\rho_{ijt}}{1-\rho_{ijt}}) = \eta_j + \Theta_i + \tau_t. \tag{17}$$

This reduced-form equation derived from the prescribing stage closely resembles our Abowd et al. (1999)-style specification in Equation (4). It differs from Equation (4) only due to the logit assumption for the error term ε.

Given its similarity to Equation (4), Equation (17) is subject to the same identification concerns. Estimated fixed effects may suffer from limited mobility bias (Andrews et al., 2012; Bonhomme et al., 2019). Moreover, logit models are known to suffer from incidental parameter bias when including fixed effects for each patient and physician. To address these issues, we follow Bonhomme et al. (2019) and Mourot (2025) and group physicians and patients in exactly the same manner as in Section 3. We then estimate the grouped parameters using the prescribing stage data via maximum likelihood. Our preferred specification uses five groups, but we demonstrate robustness of our results to the number of groups in Appendix.

Step II: Estimation of $\{\boldsymbol{\delta}, \boldsymbol{\kappa}, \bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$ from the choice stage. Conditional on our estimates for $\{\boldsymbol{\Theta}, \boldsymbol{\eta}, \boldsymbol{\tau}\}$, we estimate the parameters $\{\boldsymbol{\delta}, \boldsymbol{\kappa}, \bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$ in a second step. Recall that in Step I, we grouped both patients and physicians into five groups in our preferred specification. We maintain this grouping in Step II. In particular, since we grouped physicians

into five distinct groups, we assume that each patient differs in their information about physicians' prescribing propensity and match value only by physician group. For two physicians in the same group, a patient shares the same (mis)information and match value. Note that we do not restrict heterogeneity in information and match value across patients. As in Step I, we demonstrate robustness of our results to alternative number of groups in Appendix.

We estimate the choice stage of our model using maximum likelihood. Our approach of estimating the demand system is closely related to commonly used estimators in the literature (e.g., Goolsbee and Petrin, 2004; Ho, 2006), but additionally leverages the panel dimension of our choice stage data. Following this literature, we concentrate out the demand side mean utilities $\boldsymbol{\delta}$ by imposing a market share constraint (Berry, 1994; Berry et al., 1995). We then maximize the constrained log-likelihood of the choice stage over $\{\boldsymbol{\kappa}, \phi, \sigma_\lambda, \sigma_\zeta\}$ only. Denoting our estimates for $\{\boldsymbol{\Theta}, \boldsymbol{\eta}, \boldsymbol{\tau}\}$ from the first stage by $\widehat{\boldsymbol{\Xi}}_{\mathbf{1}}$, and defining $\boldsymbol{\Xi}_{\mathbf{2}} = \{\boldsymbol{\kappa}, \bar{\lambda}, \sigma_\lambda, \sigma_\zeta\}$, the constrained log-likelihood of the choice stage can be written as

$$\mathcal{L}_2 = \sum_{i=1}^{N} \log \int \prod_{t \in T_i} \prod_{j \in J_t} (s_{ijt}(\cdot; \boldsymbol{\Xi}_2, \widehat{\boldsymbol{\Xi}}_1))^{y_{ijt}} di \tag{18}$$
$$s.t. \quad s_{jt}(\boldsymbol{\Xi}_2, \widehat{\boldsymbol{\Xi}}_1) = \mathcal{S}_{jt},$$

where y_{ijt} is an indicator equal to 1 if patient i chooses physician j in month t, and 0 otherwise. $\mathcal{S}_{jt}$ represents the observed market share of physician j in month t in our data.

Because we use a two-step estimator, we adjust the standard errors of $\widehat{\boldsymbol{\Xi}}_2$ in the second step to account for the first-step estimation (Murphy and Topel, 1985).[23] We provide the details of these derivations in Appendix H.

Lastly, it is well known that random coefficients multinomial logit models can exhibit multiple local optima that do not correspond to the global optimum. To combat this potential problem, we follow the literature and start our optimization from 10 random starting values. Across all different versions of groupings and starting values, we find that our solver converges to the same optimum, demonstrating robustness of our results to the choice of starting values. We again provide more details on the starting values chosen in Appendix.

[23] Murphy and Topel (1985) show that under standard regularity conditions for maximum likelihood, the variance-covariance matrix for the second step, $\widehat{\boldsymbol{V}}_{\widehat{\Xi}_2}$, is given by $\widehat{\boldsymbol{V}}_{\widehat{\Xi}_2} = \widehat{\boldsymbol{V}}_2 + \widehat{\boldsymbol{V}}_2(\widehat{\boldsymbol{A}}\widehat{\boldsymbol{V}}_1\widehat{\boldsymbol{A}}')\widehat{\boldsymbol{V}}_2$, where $\widehat{\boldsymbol{V}}_1$ is the variance-covariance matrix from the first-step logit estimation, $\widehat{\boldsymbol{V}}_2$ is the variance-covariance matrix from the second-step likelihood estimation, and $\widehat{\boldsymbol{A}}$ is defined as $\widehat{\boldsymbol{A}} = \sum_i (\frac{\partial \mathcal{L}_{i2}}{\partial \widehat{\Xi}_2})(\frac{\partial \mathcal{L}_{i2}}{\partial \widehat{\Xi}_1'})$, where $\mathcal{L}_{i2}$ denotes patient i's individual contribution to the second-step likelihood.

Step III: Separating Guideline-Concordant from Guideline-Non-Concordant Patient Preferences. We decompose our estimates of patients' preference type Θ_i as in Equation (5) into three components: (i) patients' guideline-concordant preferences; (ii) observed guideline-non-concordant preferences; and (iii) a residual guideline-non-concordant preference term. We demonstrate robustness of our results to alternative number of groups in Appendix. Similarly to Step II, we adjust the standard errors of our estimates for the multi-step estimation.

6.2 Results

Table 5: Estimated Model Parameters

	Estimate	(Std. Error)		Estimate	(Std. Error)
Physician Prescribing Propensity Type η_j			**Patient Information Type Dispersion**		
η_1	-2.150	(0.038)	σ_ζ	0.168	(0.008)
η_2	-0.450	(0.021)	**Patient–Physician Match Value**		
η_4	0.279	(0.024)	σ_λ	1.623	(0.007)
η_5	1.031	(0.031)	$\bar{\lambda}$	0.619	(0.006)
Patient Preference Type Θ_i					
Θ_1	-1.529	(0.019)			
Θ_2	-0.549	(0.014)			
Θ_3	0.215	(0.014)			
Θ_4	0.876	(0.018)			
Θ_5	2.211	(0.024)			
Constant	-0.091	(0.020)			
Year FE		✓			
Data					
Patients (N)		15,019			
Physicians (J)		10,355			
Months (T)		84			
Total Observations		218,898			

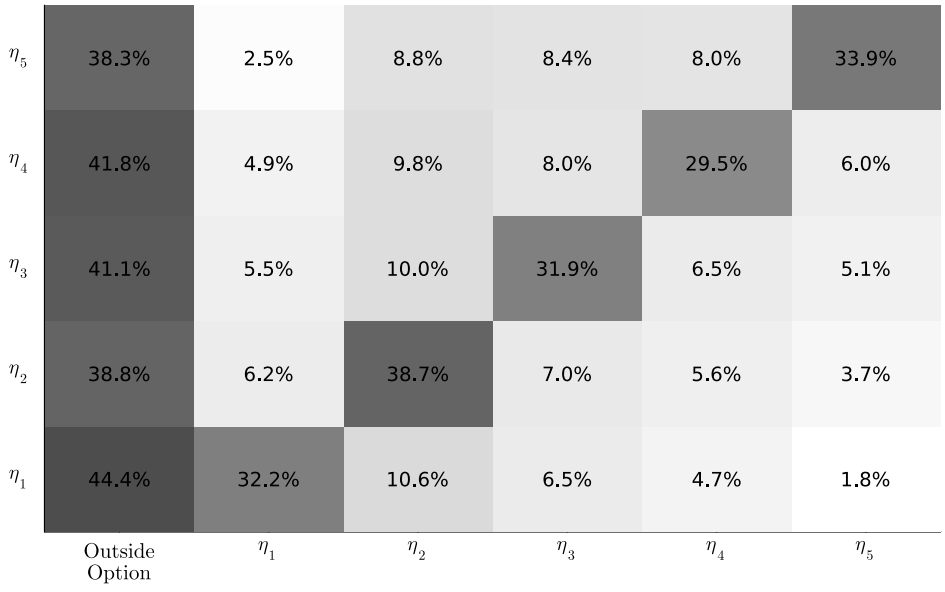
Notes: This table reports estimated model parameters from the structural model. η_3 is normalized to zero and not reported. Patient opioid preference types Θ are normalized to have weighted mean of zero. Standard errors are adjusted for two-step estimation following (Murphy and Topel, 1985) and reported in parentheses.

Table 5 reports estimates from Step 1 and Step 2 of our three-step estimation procedure using our preferred specification of five groups. In appendix , we show that the results from using 10 groups are qualitatively highly similar. In line with our motivating evidence in Section 3, our results imply that there is substantial heterogeneity in patients' preferences for opioids, $\boldsymbol{\Theta}$, and substantial dispersion in physician prescribing propensities, $\boldsymbol{\eta}$. Our results also show that patients are well aware of the differences in prescribing propensities

across physicians. The estimated standard deviation of the patient information type, σ_ζ, in Table 5 indicates that while some patients may substantially overestimate or underestimate their likelihood of obtaining prescription opioids from particular physicians during office visits, most patients appear to be well informed about physicians' prescribing propensities. Lastly, our results imply that high and low preference patients tend to choose high and low prescribing physicians for reasons other than opioid prescriptions. The estimate for the mean of the match value, $\bar{\lambda}$, implies that there is positive assortative matching even beyond an opioid prescription. However, we also find that matching based on factors other than an opioid prescription is much more dispersed, with an estimated standard deviation, σ_λ, that is about 2.5 times its mean.

Together, our estimates demonstrate that the stakes of choosing the right physician for patients are high. Consider a patient with the highest opioid preference type. Our estimates imply that the likelihood of obtaining opioids from the highest prescribing propensity physician is 90 percentage points higher than that of the lowest. These differences across physicians create strong incentives for patients to sort. Consistent with our motivating evidence in Section 3, we find strong positive sorting, owing to the relative precise information patients have about physicians' prescribing propensity. In particular, the correlation between patients' opioid preference and the chosen physicians' propensity types is 0.43.

Figure 7: Diversion ratio matrix



Notes: This figure plots the aggregate second-choice diversion ratio matrix. Each cell measures the percentage of displaced demand that switches from the origin physician group (y-axis) to the destination group (x-axis).

The strong incentives to sort are also reflected in the substitution patterns our model estimates predict. Figure 7 plots the aggregate second-choice diversion ratio matrix for our demand system, aggregated to the physician prescribing propensity group level.[24] For each of our five physician groups, these diversion ratios illustrate how much of patients' displaced demand for a given physician would divert to each physician group, or the outside option, were that physician to become unavailable. The figure shows two patterns. First, along the diagonal, patients' demand diverts disproportionately to physicians with the same prescribing propensity. Second, the largest share of patients' displaced demand flows to the outside option of forgoing care. The top row of Figure 7 illustrates both points. If a physician in group 5 were to become unavailable, 33.9% of the displaced demand would divert to a physician with the same prescribing propensity, roughly four times as much as the diversion to any other group, while 38.3% would flow to the outside option. These findings have important real-world consequences. Any policy targeting a subset of physicians for exclusion would therefore lead patients to substitute mainly to physicians with similar prescribing propensity, or to forgo care.

Lastly, we further investigate the drivers of patients' opioid preferences. Table 6 shows that patient opioid preferences correlate with both guideline-concordant and guideline-non-concordant conditions. Patients with guideline-concordant conditions, such as palliative care, cancer-related pain, and sickle cell disease, exhibit a much higher opioid preference than patients who do not suffer from these conditions. However, many guideline-non-concordant factors that appear problematic also correlate positively with opioid preferences. For instance, prior opioid misuse and substance use disorder are associated with higher opioid preferences. Additionally, we find that males and older patients have higher opioid preferences than females or younger patients.

7 Counterfactual Analysis

This section highlights the economic mechanisms through which patient sorting affects the effectiveness of top physician-targeted policies. Our counterfactuals are designed to illustrate these forces to guide policy. We first show that patient sorting attenuates policies that target top-prescribing physicians only. We then show that alternative physician-targeted policies that eliminate prescribing based on patient preferences exceeding CDC guidelines or appropriate medical consensus achieve larger and more targeted reductions in aggregate prescriptions. Appendix I provides a more detailed description of our counterfactuals and their implementation.

[24] For a more complete treatment of second-choice diversion ratios, see, e.g., Conlon and Mortimer (2021).

Table 6: Projecting Estimated Patient Opioid Preference onto Observables

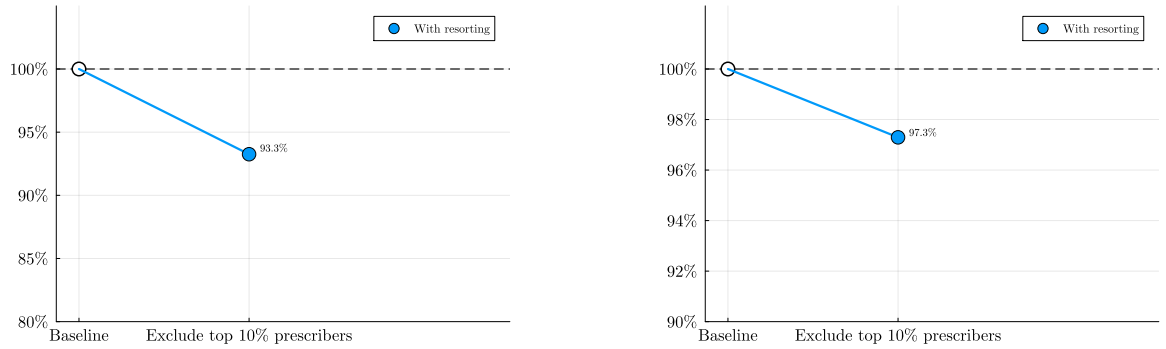
	Estimate	(Std. Error)
Conditions Concordant with Guidelines X_{1i}		
Palliative	0.186	(0.099)
Cancer-Related Pain	0.638	(0.076)
Sickle Cell Disease	0.024	(0.107)
Conditions Non-concordant with Guidelines X_{2i}		
Sleep Apnea	-0.243	(0.027)
Renal/Hepatic Deficiency	-0.107	(0.052)
Mental Health Conditions	-0.082	(0.019)
Opioid Misuse	0.722	(0.036)
Substance Use Disorder	0.074	(0.025)
Prior Nonfatal Overdose	-0.161	(0.043)
Work in Construction or Transportation	0.018	(0.035)
Female	-0.133	(0.018)
Charlson Comorbidity Index	0.000	(0.006)
Counts of Prior Inpatient Visits	0.074	(0.023)
Counts of Prior Outpatient Visits	-0.004	(0.000)
Counts of Prior ED Visits	0.017	(0.007)
Prior OOP Spending	0.000	(0.000)
Out-of-pocket Cost	0.026	(0.006)
Age 30-39	0.178	(0.040)
Age 40-49	0.164	(0.038)
Age 50-59	0.191	(0.038)
Age 60-65	0.359	(0.043)
Constant	-0.009	(0.037)
Data		
Patients (N)		15,019

Notes: Estimates from a linear regressions of $\hat{\Theta}_i$ on covariate sets X_{1i} and X_{2i}. Standard errors are adjusted for two-step estimation following Murphy and Topel (1985) and reported in parentheses. The omitted age group is below the age of 30. $R^2 = 0.079$.

7.1 Excluding Top Prescribers

First, we consider a policy that excludes the top 10% of physicians by risk-adjusted prescription rate. This policy prevents patients from choosing any of the excluded physicians. It closely mirrors real-world policies such as the DEA's prescription license revocation policy for opioid prescribing physicians, but is more extreme in the percentage of excluded physicians. We simulate counterfactual choices for all patients who saw an excluded physician in the baseline. Affected patients may either choose any of the remaining, non-excluded physicians or the outside option of not seeing any physician that month. We hold fixed the choices of patients who did not see an excluded physician and the policy functions of non-excluded physicians.[25]

Figure 8: Excluding the top 10% of physicians by risk-adjusted prescription rate



(a) Aggregate prescriptions

(b) Aggregate visits

Notes: This figure plots the aggregate prescriptions and visits after excluding the top 10% of physicians by risk-adjusted prescription rate. The numbers are expressed as percentages relative to the baseline.

We first consider the effect of the exclusion policy on aggregate prescriptions. Reducing aggregate prescriptions from the primary market is important. Opioid prescribing in the primary market is closely associated with drug overdoses and mortality, and most secondary market supply originates in the primary market (Schnell, 2025). Moreover, exposure to prescription opioids in the primary market facilitates transition into illicit opioid use (Eichmeyer and Zhang, 2022).[26] The left panel of Figure 8 compares the counterfactual aggregate prescriptions in percentages relative to the baseline. Perhaps surprisingly, our policy reduces

[25] Currie et al. (2023) finds that physicians respond to competitive pressure from nurse practitioners by prescribing more liberally. Modeling competition is outside the scope of this paper, and we hold fixed physicians' policy functions, abstracting away from physician competition.

[26] Reducing prescriptions in the primary market does not come without costs. Patients who currently rely on prescriptions from the primary market may substitute to the secondary market in the short run. However, lowering opioid supply in the primary market may substantially reduce illicit opioid use and opioid misuse in the long run (Eichmeyer and Zhang, 2022; Schnell, 2025).

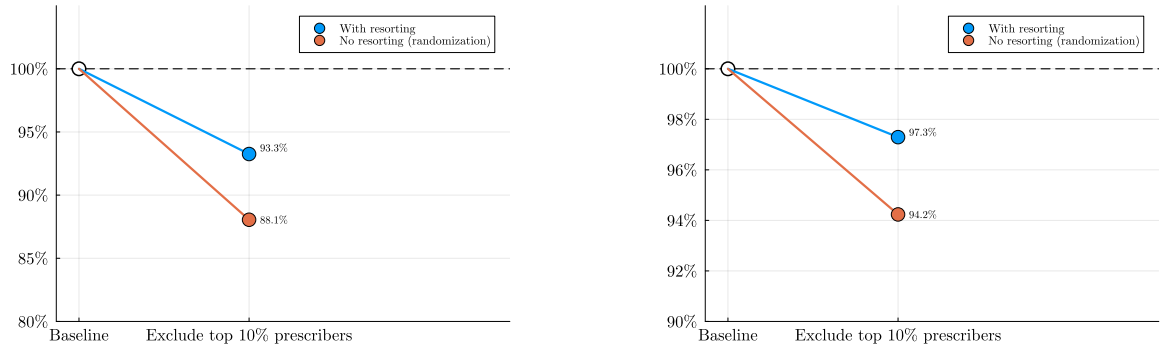
aggregate prescriptions by only about 7 percentage points. Given that the top 10% of physicians account for 32% of aggregate prescriptions, all else equal, excluding high-prescribing physicians should have reduced prescriptions more dramatically.[27]

Next, aggregate visits decline under the exclusion policy as illustrated in the right panel of Figure 8. This implies that a fraction of the observed reduction in the aggregate prescriptions is driven by patients forgoing care. We further decompose the reduction in aggregate prescriptions into the contribution of affected patients switching to lower-prescribing physicians and patients forgoing care.[28] Our results imply that 74% of the predicted reduction in aggregate prescriptions is attributable to patients being forced to resort to physicians with lower prescribing propensities, while the remaining 26% is caused by patients forgoing care.

7.2 Quantifying Patient Resorting

To quantify the importance of the resorting effect, we simulate an exclusion policy but explicitly prohibit resorting. In particular, affected patients can no longer choose their physician. Instead, affected patients only choose between being randomly assigned to one of the remaining physicians and not seeking care that month. We hold fixed the choices of patients who did not see an excluded physician and the policy functions of non-excluded physicians.

Figure 9: Excluding the top 10% of physicians and banning patient resorting



(a) Aggregate prescriptions

(b) Aggregate visits

Notes: This figure plots the aggregate prescriptions and visits after excluding the top 10% of physicians by risk-adjusted prescription rate and randomly assigning patients to the remaining physicians. The numbers are expressed as percentages relative to the baseline.

[27] If affected patients switched randomly to the remaining physicians, a simple back-of-the-envelope calculation, without the structural model, would suggest a reduction of about 11 percentage points.

[28] To implement this decomposition, we simulate the reduction in prescriptions by requiring each patient to stay in care. The resulting reduction isolates the contribution of patients switching to lower-prescribing physicians. We then compare the resulting reduction to that when we do not require patients to stay in care.

The left panel of Figure 9 compares aggregate prescriptions resulting from this counterfactual policy to the exclusion policy of Section 7.1 relative to the baseline. When patients cannot reoptimize their choice of physician, aggregate opioid prescriptions drop by an additional 5 percentage points. Therefore, patient resorting substantially attenuates the effect of our exclusion policy. The right panel of Figure 9 shows the aggregate visits compared to the exclusion policy of Section 7.1 relative to the baseline. Aggregate visits drop by an additional 3 percentage points when patients are not allowed to resort. In other words, the larger reduction in aggregate opioid prescriptions comes at the cost of many patients forgoing care.

The additional reduction in aggregate prescriptions originates from two channels. First, patients who remain in the market can no longer resort to physicians whose prescribing propensities are, ceteris paribus, close to those of the excluded physicians. Second, prohibiting resorting leads some patients to forgo care altogether. We decompose the additional 5 percentage points reduction in aggregate prescriptions into these two channels. We find that 8% of the reduction is due to the elimination of patient resorting, while the remaining 92% is driven by patients forgoing care. These results imply that a non-trivial fraction of patients who unequivocally qualify for opioid treatment under standard clinical guidelines or medical consensus are prevented from obtaining appropriate care. This constitutes a clearly undesirable policy outcome.

7.3 Breaking Bad Opioid Sorting

The previous subsection highlights a trade-off between curbing opioid prescribing and the risk of denying access to clinically appropriate care. In this subsection, we explicitly address this trade-off. We consider alternative policies that aim to preserve opioid prescribing supported by CDC guidelines or appropriate medical consensus, while eliminating prescribing based on preferences concordant with neither. Defining appropriate medical consensus in our setting is challenging for two reasons. First, the CDC guidelines define only two diagnoses as unequivocally guideline-concordant. Prescribing based on any other medical consensus, such as prescribing based on pain level, is guideline-non-concordant but may still be considered appropriate by some.[29] Second, medical consensus is inherently unobservable in claims data. To overcome these challenges, we bound appropriate medical consensus using two alternative extreme assumptions:

(A) Any prescribing based on preferences exceeding the CDC guidelines is inappropriate.

[29] How much consensus-based prescribing is appropriate is itself an open question. However, there is growing evidence that opioid prescribing for pain as part of the medical consensus may have contributed to drug deaths (Evans and Lieber, 2026).

(B) Any prescribing based on preferences the CDC guidelines do not caution against counts as appropriate medical consensus.

In our baseline, physicians prescribe based on all patient preferences. Consequently, patients sort based on both guideline-concordant and guideline-non-concordant preferences. We categorize sorting driven by patient preferences concordant with CDC guidelines or any potential appropriate medical consensus as "good sorting" and the remainder as "bad sorting." In our counterfactual simulations, we break "bad" sorting while preserving "good" sorting by requiring physicians to prescribe based only on preferences concordant with CDC guidelines or appropriate medical consensus. As before, we hold fixed the choices of unaffected patients.

Specifically, for each patient affected by the exclusion policy, we require physicians to prescribe based on the following rules:

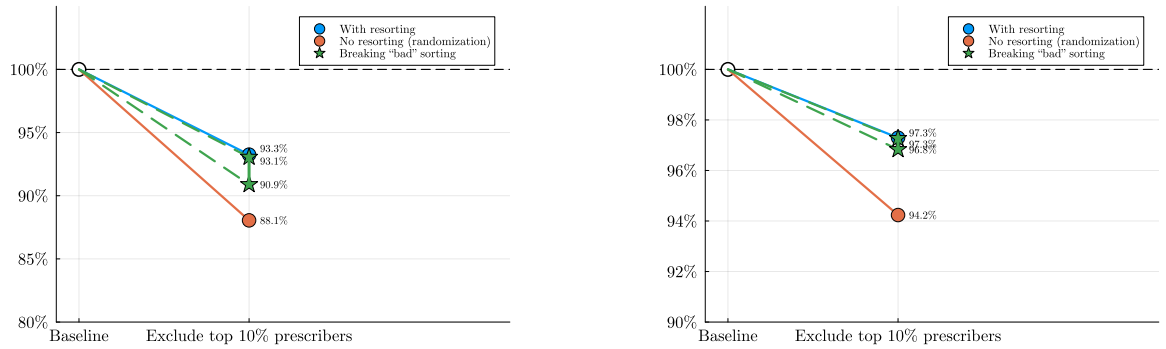
1. If the patient has cancer-related pain or sickle cell disease, a physician considers the full opioid preference type Θ. In other words, we do not interfere in opioid prescribing for chronic pain concordant with the CDC guidelines.

2. In all other cases, the physician disregards preferences exceeding the guidelines or appropriate medical consensus.

The left panel of Figure 10 compares the reductions in aggregate prescriptions under the breaking "bad" sorting exclusion policy for both assumptions A and B to the exclusion policies with and without resorting. Breaking "bad" sorting reduces aggregate prescriptions by 7 to 9 percentage points. The upper bound of our range corresponds to the policy effect under assumption A. Under this assumption, the magnitude of the reduction in aggregate prescriptions is comparable to that of the exclusion policy without resorting (12 percentage points). The lower bound of our range corresponds to the policy effect under assumption B. Even under this most generous interpretation of appropriate medical consensus, our counterfactual policy results in a meaningful additional reduction in aggregate prescriptions compared to the exclusion policy with resorting.

Our results indicate that the breaking "bad" sorting exclusion policy can reduce aggregate prescriptions more effectively than the exclusion policy with resorting. However, the policy does so while keeping patients in care. The right panel of Figure 10 compares aggregate visits under the breaking "bad" sorting exclusion policy for both assumptions A and B to the exclusion policies with and without resorting. Regardless of our assumptions on appropriate medical consensus, resulting office visits closely match that of the exclusion policy with

Figure 10: Excluding the top 10% of physicians and breaking sorting based on preferences concordant with neither guidelines nor appropriate medical consensus



(a) Aggregate prescriptions

(b) Aggregate visits

Notes: This figure plots the aggregate prescriptions and visits after excluding the top 10% of physicians by risk-adjusted prescription rate and eliminating sorting based on preferences concordant with neither guidelines nor appropriate medical consensus, under different assumptions on appropriate medical consensus. The numbers are expressed as percentages relative to the baseline.

resorting. Appendix Figure K1 further demonstrates that these results do not depend on the choice of cutoff value used to exclude physicians.

Finally, we consider a counterfactual policy that does not rely on the exclusion of physicians. Exclusionary policies might be hard to implement at scale in reality because they involve significant legal risk. Instead, we consider the breaking "bad" sorting policy for all patients without excluding any physician from the market. This policy completely eliminates discretionary prescribing for all physicians. The breaking "bad" sorting policy may be feasible if regulators or insurers can identify guideline-concordant conditions and appropriate medical consensus to audit physicians' prescribing. We find that this policy reduces aggregate prescriptions by up to 29 percentage points, depending on the assumption on appropriate medical consensus. The magnitude of this result is comparable to an exclusion policy with resorting that excludes the top 15% to 18% of physicians by risk-adjusted prescription rate. Simultaneously, breaking "bad" sorting preserves access to care for patients whose preferences are consistent with CDC guidelines and appropriate medical consensus. In particular, aggregate visits drop by at most 4 percentage points.

8 Conclusion

In this paper, we provide novel evidence on how consumers sort into experts. We demonstrate that patients with chronic pain strongly sort into physicians based on their preferences for prescription opioids. We then estimate a structural model of patient sorting and physician

prescribing to understand the policy implications of patient sorting.

Through counterfactual analyses, we quantify how sorting affects policies targeting high-prescribing physicians, thereby substantially attenuating their effectiveness. Moreover, nearly all of this attenuation of top physician-targeted policies arises from resorting driven by patient preferences concordant with neither clinical guidelines nor appropriate medical consensus. We further show that relatively straightforward policies that eliminate physician discretion in prescribing can eliminate undesirable patient sorting while still preserving patients' necessary clinical needs.

Although we study opioid prescribing, the sorting mechanism applies in other expert markets where outcomes are driven by consumers sorting into willing experts. In the financial advice market, firms bar individual advisers for misconduct, but this discipline is partly undone as unsophisticated consumers sort into other misconduct-prone advisers (Egan et al., 2019). In the consumer lending market, regulators ban individual predatory lenders from the industry (Federal Trade Commission, 2021), but borrowers may resort to other willing lenders. As in our setting, supply-side policies that target individual experts may be improved by leveraging consumer resorting.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis (1999). High wage workers and high wage firms. *Econometrica 67*(2), 251–333.

Alpert, A., S. Dykstra, and M. Jacobson (2024). Hassle costs versus information: How do prescription drug monitoring programs reduce opioid prescribing? *American Economic Journal: Economic Policy 16*(1), 87–123.

Alpert, A., W. N. Evans, E. M. J. Lieber, and D. Powell (2021). Origins of the opioid crisis and its enduring impacts. *The Quarterly Journal of Economics 137*(2), 1139–1179.

Alpert, A., D. Powell, and R. L. Pacula (2018). Supply-side drug policy in the presence of substitutes: Evidence from the introduction of abuse-deterrent opioids. *American Economic Journal: Economic Policy 10*(4), 1–35.

Alpert, A., S. D. Schwab, and B. Ukert (2025). Opioid use and employment outcomes: Evidence from the u.s. military. *Review of Economics and Statistics*, 1–45.

Andrews, M., L. Gill, T. Schank, and R. Upward (2012). High wage workers match with high wage firms: Clear evidence of the effects of limited mobility bias. *Economics Letters 117*(3), 824–827.

Badinski, I., A. Finkelstein, M. Gentzkow, and P. Hull (2023). Geographic variation in healthcare utilization: The role of physicians. Working Paper 31749, National Bureau of Economic Research.

Baker, D. W. (2017). History of the joint commission's pain standards: lessons for today's pre-

scription opioid epidemic. *Jama 317*(11), 1117–1118.

Barnett, M. L., A. R. Olenski, and A. B. Jena (2017). Opioid-prescribing patterns of emergency physicians and risk of long-term use. *New England Journal of Medicine 376*(7), 663–673.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics 25*(2), 242–262.

Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica 63*(4), 841–890.

Bonhomme, S., K. Holzheu, T. Lamadon, E. Manresa, M. Mogstad, and B. Setzler (2023). How much should we trust estimates of firm effects and worker sorting? *Journal of Labor Economics 41*(2), 291—-322.

Bonhomme, S., T. Lamadon, and E. Manresa (2019). A distributional framework for matched employer employee data. *Econometrica 87*(3), 699–739.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics 36*(S1), S13–S70.

Carpenter, C. S., C. B. McClellan, and D. I. Rees (2017). Economic conditions, illicit drug use, and substance use disorders in the united states. *Journal of health economics 52*, 63–73.

Case, A. and A. Deaton (2021). Deaths of despair and the future of capitalism.

Chandra, A., A. Finkelstein, A. Sacarny, and C. Syverson (2016). Health care exceptionalism? performance and allocation in the us health care sector. *American Economic Review 106*(8), 2110–44.

Charlson, M. E., P. Pompei, K. L. Ales, and C. MacKenzie (1987). A new method of classifying prognostic comorbidity in longitudinal studies: Development and validation. *Journal of Chronic Diseases 40*(5), 373–383.

Chou, R., D. Hartung, J. Turner, I. Blazina, B. Chan, X. Levande, M. McDonagh, S. Selph, R. Fu, and M. Pappas (2020). Opioid treatments for chronic pain. *Rockville (MD): Agency for Healthcare Research and Quality (US); (Comparative Effectiveness Review, No. 229.)*.

Conlon, C. and J. H. Mortimer (2021, December). Empirical properties of diversion ratios. *The RAND Journal of Economics 52*(4), 693–726.

Currie, J., A. Li, and M. Schnell (2023). The effects of competition on physician prescribing. Technical report, National Bureau of Economic Research.

Currie, J. M. and W. B. MacLeod (2020). Understanding doctor decision making: The case of depression treatment. *Econometrica 88*(3), 847–878.

Dahlhamer, J. M., E. M. Connor, J. Bose, J. W. Lucas, and C. E. Zelaya (2021). Prescription opioid use among adults with chronic pain: United states, 2019. *National Health Statistics Reports; no 162. Hyattsville, MD: National Center for Health Statistics. 2021.*.

Deyo, R. (1992). Adapting a clinical comorbidity index for use with icd-9-cm administrative databases. *Journal of Clinical Epidemiology 45*(6), 613–619.

Ding, Y. and C. Liu (2021). Alternative payment models and physician treatment decisions: Evidence from lower back pain. *Journal of Health Economics 80*, 102548.

Dowell, D., T. M. Haegerich, and R. Chou (2016). CDC Guideline for Prescribing Opioids for Chronic Pain-United States, 2016. *MMWR-Morbidity and Mortality Weekly Report, 65*.

Dowell, D., K. R. Ragan, C. M. Jones, G. T. Baldwin, and R. Chou (2022). CDC Clinical Practice Guideline for Prescribing Opioids for Pain—United States, 2022. *MMWR. Recommendations and Reports 71*(3), 1–95.

Egan, M., G. Matvos, and A. Seru (2019). The market for financial adviser misconduct. *Journal of Political Economy 127*(1), 233–295.

Eichmeyer, S. and J. Zhang (2022). Pathways into opioid dependence: Evidence from practice variation in emergency departments. *American Economic Journal: Applied Economics 14*(4), 271–300.

Evans, W. N. and E. M. Lieber (2026, January). Prescription for disaster: The ssdi rate, pain, and prescribing practices. *Journal of Human Resources*, 0724–13685R2.

Evans, W. N., E. M. Lieber, and P. Power (2019). How the reformulation of oxycontin ignited the heroin epidemic. *Review of Economics and Statistics 101*(1), 1–15.

Federal Trade Commission (2021). FTC acts to ban payday lender from industry, forgive illegal debt. Press release, Federal Trade Commission. In the Matter of Harvest Moon Financial, Gentle Breeze Online, and Green Stream Lending.

Finkelstein, A., M. Gentzkow, and D. Li (2025). What drives risky prescription opioid use? evidence from migration. *The Quarterly Journal of Economics 140*(4), 3133–3189.

Finkelstein, A., M. Gentzkow, and H. Williams (2016). Sources of geographic variation in health care: Evidence from patient migration*. *The Quarterly Journal of Economics 131*(4), 1681–1726.

Goolsbee, A. and A. Petrin (2004). The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica 72*(2), 351–381.

Gui, L., C. Qin, and M. Xiao (2024). Anatomy of opioid diversion: Examining supply-side curtailment. *Available at SSRN 5044557*.

Ho, K. (2006). The welfare effects of restricted hospital choice in the us medical care market. *Journal of Applied Econometrics 21*(7), 1039–1079.

Hollingsworth, A., C. J. Ruhm, and K. Simon (2017). Macroeconomic conditions and opioid abuse. *Journal of health economics 56*, 222–233.

Hsu, J., J. Schmittdiel, E. Krupat, T. Stein, D. Thom, B. Fireman, and J. Selby (2003). Patient choice. a randomized controlled trial of provider selection. *Journal of General Internal Medicine 18*, 319–325.

Janssen, A. and X. Zhang (2023). Retail pharmacies and drug diversion during the opioid epidemic. *American economic review 113*(1), 1–33.

Kline, P. (2024). Chapter 2 - firm wage effects. Volume 5 of *Handbook of Labor Economics*, pp. 115–181. Elsevier.

Lee, M., S. M. Silverman, H. Hansen, V. B. Patel, and L. Manchikanti (2011). A comprehensive review of opioid-induced hyperalgesia. *Pain physician 14*(2), 145.

Lucas, J. W. and I. Sohi (2024). Chronic pain and high-impact chronic pain in u.s. adults, 2023. *NCHS data brief, no 518, Hyattsville, MD: National Center for Health Statistics. 2024.*.

Meara, E., J. R. Horwitz, W. Powell, L. McClelland, W. Zhou, A. J. O'Malley, and N. E. Morden (2016). State legal restrictions and prescription-opioid use among disabled adults. *New England Journal of Medicine 375*(1), 44–53.

Molitor, D. (2018). The evolution of physician practice styles: Evidence from cardiologist migration. *American Economic Journal: Economic Policy 10*(1), 326–56.

Mourot, P. (2025). Should top surgeons practice at top hospitals? sorting and complementarities in healthcare. Working paper, Boston University.

Mularski, R. A., F. White-Chu, D. Overbay, L. Miller, S. M. Asch, and L. Ganzini (2006). Measuring pain as the 5th vital sign does not improve quality of pain management. *Journal of general internal medicine 21*(6), 607–612.

Murphy, K. M. and R. H. Topel (1985). Estimation and inference in two-step econometric models. *Journal of Business Economic Statistics 3*(4), 370–379.

Office of the Inspector General (2019). Review of the drug enforcement administration's regulatory and enforcement efforts to control the diversion of opioids. Technical report, U.S. Department of Justice.

Sacarny, A., T. Avilova, I. Williamson, W. Merrick, and M. Jacobson (2025). Prescription drug monitoring program reminder emails, program use, and prescribing: A randomized clinical trial. In *JAMA Health Forum*, Volume 6, pp. e255623. American Medical Association.

Sacarny, A., D. Yokum, A. Finkelstein, and S. Agrawal (2016). Medicare letters to curb overprescribing of controlled substances had no detectable effect on providers. *Health Affairs 35*(3), 471–479.

Schneeweiss, S., P. S. Wang, J. Avorn, and R. J. Glynn (2003). Improved comorbidity adjustment for predicting mortality in medicare populations. *Health Services Research 38*(4), 1103–1120.

Schnell, M. (2025). Physician behavior in the presence of a secondary market: The case of prescription opioids. Econometrica, forthcoming.

Schnell, M. and J. Currie (2018). Addressing the opioid epidemic: Is there a role for physician education? *American Journal of Health Economics 4*(3), 383–410.

Soliman, A. (2025). Disrupting drug markets: The effects of crackdowns on rogue opioid suppliers.

American Economic Journal: Economic Policy 17(4), 165–191.

Staiger, B., L. Baker, and T. Hernandez-Boussard (2022). Provider opioid prescribing behaviors and opioid use in medicaid. Working paper.

Volkow, N. D. and T. McLellan (2016). Opioid abuse in chronic pain–misconceptions and mitigation strategies. *The New England Journal of Medicine 374*, 1253–1263.

Online Appendix

A CDC Clinical Practice Guideline for Prescribing Opioids for Pain

CDC2016 and CDC2022 offer recommendations for clinicians who provide pain care, including opioid prescribing, to outpatients aged 18 or above. Following these guidelines, we classify medical conditions as either concordant (Table A1) or non-concordant (Table A2) with the guidelines. We label a condition concordant when the guidelines place it outside the scope of their chronic-pain recommendations, declining to advise against opioid treatment for it (cancer-related pain, palliative or end-of-life care, and sickle cell disease).[30] We label a condition non-concordant when the guidelines recommend against opioids or urge caution in prescribing them.

Medical Diagnosis	**ICD10 Diagnosis Codes**
sickle cell disease	D57
cancer-related pain	G89.3
palliative / end-of-life care	Z51.5

Table A1: Medical conditions that are concordant with CDC guidelines

Medical Diagnosis	**ICD10 Diagnosis Codes**
sleep-disordered breathing	G47.3
renal or hepatic insufficiency	N17, K72
mental health condition (anxiety, depression and post-traumatic stress disorder)	F32, F41, F43
opioid abuse	F11, T40
substance use disorder	F10, F12-F19
non-fatal overdose	T36 - T39, T41 - T50

Table A2: Medical conditions that are non-concordant with CDC guidelines

[30] Cancer-related pain and palliative or end-of-life care fall into this category under both guidelines. Sickle cell disease additionally falls into this category under CDC2022.

B Charlson Comorbidity Index (CCI)

We use weights in Schneeweiss et al. (2003) for each of the 17 conditions suggested for the Charlson/Deyo variant (Deyo, 1992). The 17 conditions are:

Comorbid Condition	**ICD10 Diagnosis Codes**	**Weights**
Myocardial Infarction	I21, I22, I25.2	1
Congestive Heart Failure	I09.9, I11.0, I13.0, I13.2, I25.5, I42.0, I42.5–I42.9, I43, I50, P29.0	1
Peripheral Vascular Disease	I70, I71, I73.1, I73.8, I73.9, I77.1, I79.0, I79.2, K55.1, K55.8, K55.9, Z95.8, Z95.9	1
Cerebrovascular Disease	G45, G46, H34.0, I60–I69	1
Dementia	F00–F03, F05.1, G30, G31.1	1
Chronic Pulmonary Disease	I27.8, I27.9, J40–J47, J60–J67, J68.4, J70.1, J70.3	1
Connective Tissue Disease	M05, M06, M31.5, M32–M34, M35.1, M35.3, M36.0	1
Peptic Ulcer Disease	K25–K28	1
Mild Liver Disease	B18, K70.0–K70.3, K70.9, K71.3–K71.5, K71.7, K73, K74, K76.0, K76.2–K76.4, K76.8, K76.9, Z94.4	1
Diabetes w/o Complications	E10.0–E10.1, E10.6, E10.8–E10.9, E11.0–E11.1, E11.6, E11.8–E11.9, E12.0–E12.1, E12.6, E12.8–E12.9, E13.0–E13.1, E13.6, E13.8–E13.9, E14.0–E14.1, E14.6, E14.8–E14.9	1
Diabetes w/ Chronic Complications	E10.2–E10.5, E10.7, E11.2–E11.5, E11.7, E12.2–E12.5, E12.7, E13.2–E13.5, E13.7, E14.2–E14.5, E14.7	2
Paraplegia and Hemiplegia	G04.1, G11.4, G80.1, G80.2, G81, G82, G83.0–G83.4, G83.9	2
Renal Disease	I12.0, I13.1, N03.2–N03.7, N05.2–N05.7, N18, N19, N25.0, Z49.0–Z49.2, Z94.0, Z99.2	2
Cancer	C00–C26, C30–C34, C37–C41, C43, C45–C58, C60–C76, C81–C85, C88, C90–C97	2
Moderate or Severe Liver Disease	I85.0, I85.9, I86.4, I98.2, K70.4, K71.1, K72.1, K72.9, K76.5–K76.7	3
Metastatic Carcinoma	C77–C80	6
HIV/AIDS	B20–B22, B24	6

Table B1: Charlson Comorbidity Index: ICD10 Diagnoses and Weights

Patients in the following age groups receive additional points for their index:

Age Groups	**Weights**
$<$50 years	0
50–59 years	1
60–69 years	2
70–79 years	3
$\geq$80 years	4

Table B2: Charlson Comorbidity Index: Age Groups and Weights

C Constructing Expected Out-of-Pocket Costs

In our data, out-of-pocket (OOP) opioid costs are observed only for chronic pain patients who fill an opioid prescription. We therefore use expected out-of-pocket cost per day. It is the average OOP opioid cost per day at the plan–year–month level among observed users.

$$\bar{OOP}_{it} = \frac{1}{N_{plan(i)t}} \sum_{i' \in \text{plan}(i)t} OOP_{i't}.$$

$N_{plan(i)t}$ is the number of patients in i's plan in year-month t who filled an opioid prescription.

D Grouping Algorithm

To group both patients and providers, we first run two regressions to construct the average risk-adjusted prescription rate across patients for providers, and the average risk-adjusted prescription rate across providers for patients. In particular, we run

$$\text{prescribe}_{ijt} - X_{it}\beta = \tilde{\eta}_j + \epsilon_{ijt}, \tag{D1}$$

and

$$\text{prescribe}_{ijt} - X_{it}\beta = \tilde{\theta}_i + \epsilon_{ijt}, \tag{D2}$$

where X_{it} is the patient characteristics defined in Section 3. We then use our estimates for $\tilde{\eta}_j$ and $\tilde{\theta}_i$ to group patients and providers using the k-mean clustering algorithm (Bonhomme et al., 2019) as in Mourot (2025). In the second stage, we can then consistently estimate the two-way fixed effects model in Equation (4) using our groups from this grouping stage.

E Robustness: Variance Decomposition of Prescription Decision and Opioid Misuse (%)

Table E1: Robustness: Variance Decomposition of Prescription Decision (%)

	Baseline	**Alt. Sample Restrictions,** $K = 5$				**Bayesian Shrinkage**	**Including** X_{it}
	$K = 5$ ≥ 2 visits/patient ≥ 2 visits/physician	≥ 1 visits/patient ≥ 1 visits/physician	≥ 5 visits/patient ≥ 5 visits/physician	≥ 10 visits/patient ≥ 10 visits/physician	Including patients who never received an opioid prescription	$K = 5$ Bayes shrinkage	$K = 5$ X_{it} included
Patient							
$\frac{\text{Var}(X_{2i}\beta_2 + \theta_i)}{\text{Var}(\text{prescribe}^E)}$	52.36	48.44	59.77	63.97	82.32	52.95	52.23
$\frac{\text{Var}(X_{1i}\beta_1)}{\text{Var}(\text{prescribe}^E)}$	0.05	0.05	0.05	0.03	0.20	2.69	0.05
Physician							
$\frac{\text{Var}(\eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$	14.78	17.99	9.67	7.27	2.00	12.51	14.83
Sorting							
$2 \times \frac{\text{Cov}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$	32.98	33.67	30.68	28.85	15.25	30.39	33.05
$2 \times \frac{\text{Cov}(X_{1i}\beta_1, \eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$	-0.03	-0.02	-0.03	-0.01	0.04	1.44	-0.03
Correlation							
$\text{corr}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})$	0.59	0.57	0.64	0.67	0.59	0.59	0.59
$\text{corr}(X_{1i}\beta_1, \eta_{j(i,t)})$	-0.02	-0.01	-0.02	-0.01	0.03	0.12	-0.02

Notes: This table decomposes the variance in the prescription decision into patient, physician, and sorting components, with each component scaled by $\text{Var}(\text{prescribe}^E_{it})$, the variation in prescribing explained by patient- and physician-side factors. All columns use $K = 5$ groups for patients and physicians. The last column of the alternative sample restrictions additionally includes patients who never received an opioid prescription during the sample period. The "Bayesian Shrinkage" column applies an empirical Bayes-adjusted estimator to the patient and physician fixed effects used as the input for grouping via k-means clustering. The "Including X_{it}" column adds time-varying patient characteristics X_{it} in Equation (4) and treats X_{it} as guideline-non-concordant.

Table E2: Robustness: Variance Decomposition of Opioid Misuse (%)

	Baseline	**Alt. Sample Restrictions,** $K = 5$				**Bayesian Shrinkage**	**Including** X_{it}
	$K = 5$ ≥ 2 visits/patient ≥ 2 visits/physician	≥ 1 visits/patient ≥ 1 visits/physician	≥ 5 visits/patient ≥ 5 visits/physician	≥ 10 visits/patient ≥ 10 visits/physician	Including patients who never received an opioid prescription	$K = 5$ Bayes shrinkage	$K = 5$ X_{it} included
Patient							
$\frac{\text{Var}(X_{2i}\beta_2 + \theta_i)}{\text{Var}(\text{misuse}^E)}$	75.43	72.03	74.61	81.98	86.53	81.43	75.41
$\frac{\text{Var}(X_{1i}\beta_1)}{\text{Var}(\text{misuse}^E)}$	0.05	0.04	0.07	0.08	0.07	0.70	0.05
Physician							
$\frac{\text{Var}(\eta_{j(i,t)})}{\text{Var}(\text{misuse}^E)}$	4.67	6.09	5.07	2.47	1.62	2.95	4.68
Sorting							
$2 \times \frac{\text{Cov}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})}{\text{Var}(\text{misuse}^E)}$	19.87	21.86	20.30	15.50	11.71	14.71	19.88
$2 \times \frac{\text{Cov}(X_{1i}\beta_1, \eta_{j(i,t)})}{\text{Var}(\text{misuse}^E)}$	0.00	0.00	-0.00	-0.00	0.02	0.17	0.00
Correlation							
$\text{corr}(X_{2i}\beta_2 + \theta_i, \eta_{j(i,t)})$	0.53	0.52	0.52	0.55	0.49	0.47	0.53
$\text{corr}(X_{1i}\beta_1, \eta_{j(i,t)})$	0.00	0.00	-0.00	-0.00	0.03	0.06	0.00

Notes: This table decomposes the variance in opioid misuse into patient, physician, and sorting components, with each component scaled by $\text{Var}(\text{misuse}^E_{it})$, the variation in opioid misuse explained by patient- and physician-side factors. All columns use $K = 5$ groups for patients and physicians. The last column of the alternative sample restrictions additionally includes patients who never received an opioid prescription during the sample period. The "Bayesian Shrinkage" column applies an empirical Bayes-adjusted estimator to the patient and physician fixed effects used as the input for grouping via k-means clustering. The "Including X_{it}" column adds time-varying patient characteristics X_{it} in Equation (4) and treats X_{it} as guideline-non-concordant.

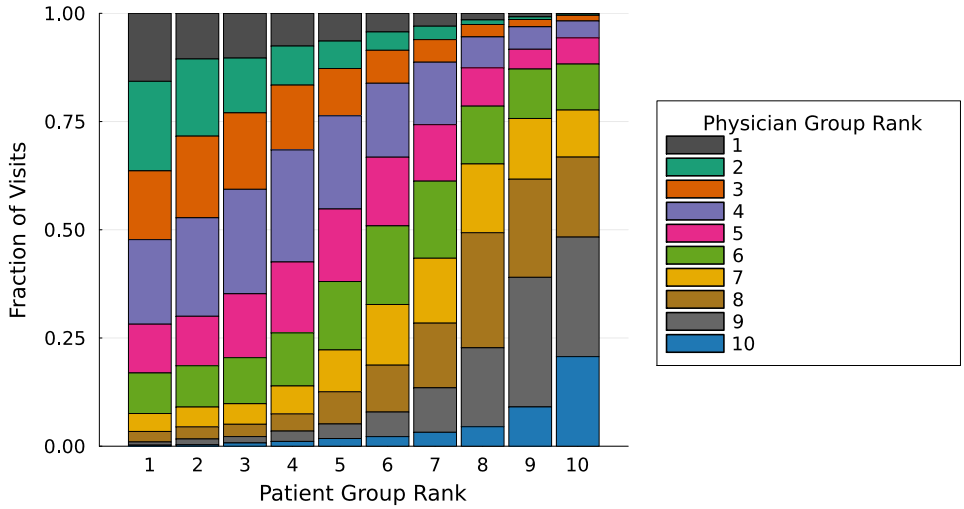
Table E3: Variance Decomposition of Prescription Decision Using Subsamples (%)

	Restrictions on Number of Prior Visits			
	no restriction (baseline)	≥ 1	≥ 2	≥ 3
Sorting				
$2 \times \dfrac{\text{Cov}(\Theta_i, \eta_{j(i,t)})}{\text{Var}(\text{prescribe}^E)}$	32.95	28.19	26.16	25.45

Notes: This table reports the sorting component of the variance decomposition of the prescription decision, scaled by $\text{Var}(\text{prescribe}_{it}^E) \equiv \text{Var}(\Theta_i + \eta_{j(i,t)})$, the variation in prescribing explained by patient- and physician-side factors. Each column uses a different subsample. These subsamples differ in the minimum number of prior visits per patient-physician pair.

F Patient Sorting

Figure F1: Patient Sorting Using 10 Groups



Notes: This figure plots, for each of the 10 patient groups, the distribution of visits across the 10 physician groups. Patient and physician groups are formed from their respective fixed effects estimates and ranked from group 1 (lowest type) to group 10 (highest type).

G Constructing Opioid Misuse Measure

Following the literature, we classify an outpatient visit as involving opioid misuse if the patient's daily Morphine Equivalent Dosage (MED) exceeds 120 mg, given the documented correlation between such prescription levels and opioid abuse (Meara et al., 2016; Staiger

et al., 2022; Finkelstein et al., 2025). In this section, we describe how we calculate the MED for an outpatient visit.

For a visit t, we observe all prescription drugs d, their active opioid ingredients o, and the respective dosages in milligrams (mg).[31] Following the CDC guidelines (CDC2016), we consider commonly prescribed opioid ingredients in pain management: codeine, fentanyl, hydrocodone, hydromorphone, methadone, morphine, oxycodone, oxymorphone, tapentadol, and tramadol. For each prescription drug d, we define its morphine equivalent dosage (MED) as the sum of all opioids' dosage multiplied by the conversion factor:

$$MED_d = \sum_{o \in N_d^{opioid}} dosage_o \times conversion_o$$

N_d^{opioid} denotes all the opioid ingredients found in drug d. $conversion_o$ is the conversion factor used by the CDC to standardize the comparisons between opioids relative to morphine and is found in (CDC2016). Hence, patient i's daily MED in outpatient visit t is defined as the sum of all drugs' total daily MED:

$$MED_{it} = \sum_{d \in N_{it}^{drug}} \frac{MED_d \times METQTY_{dit}}{DAYSUPP_{dit}}$$

N_{it}^{drug} denotes all the drugs prescribed to patient i in outpatient visit t. $METQTY_{dit}$ is the metric quantity of drug d for patient i in outpatient visit t (i.e., number of pills prescribed). $DAYSUPP_{dit}$ is drug d's days of supply for patient i in outpatient visit t. Hence, an outpatient visit t is defined as an opioid misuse if MED_{it} exceeds 120 mg.

H Gradient for Likelihood Equation (18)

In this section, we describe how we derive the gradient for the likelihood Equation (18). We use these derivations to calculate the Murphy-Topel adjusted standard errors in Equations (**??**) and (**??**).

We can now derive the following partial derivatives. For each element m of the ϕ vector and with some abuse of notation (integral vs sum), we have

$$\frac{\partial s_{jt}}{\partial \phi^{\{m\}}} = \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \phi^{\{m\}}} \, dF(\zeta) = \sum_i \int_{-\infty}^{+\infty} z_i^{\{m\}} s_{ijt} s_{i0t} \, dF(\zeta). \tag{H1}$$

[31] Dosages in micrograms (mcg) are divided by 1,000 to convert to milligrams (mg).

ζ random coefficient σ_ζ

The parameter σ_ζ enters through two channels: the prescription probability $\tilde{\rho}_{ijt}$ and the interaction term $\Theta_i\omega_{ij}$. Thus we have:

$$\begin{aligned}\frac{\partial s_{jt}}{\partial \sigma_\zeta} &= \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \sigma_\zeta}\, dF(\zeta) \\ &= \sum_i \int_{-\infty}^{+\infty} s_{ijt}(\Theta_i + \alpha X_{it}) \left[\zeta_j(\tilde{\rho}_{ijth}(1-\tilde{\rho}_{ijth}) + \psi) - \sum_l \zeta_j(\tilde{\rho}_{ilth}(1-\tilde{\rho}_{ilth}) + \psi) s_{ilth}\right] dF(\zeta).\end{aligned} \tag{H2}$$

λ random coefficient σ_λ

The parameter σ_λ enters utility additively through $\sigma_\lambda\lambda_j$

$$\begin{aligned}\frac{\partial s_{jt}}{\partial \sigma_\lambda} &= \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \sigma_\lambda}\, dF(\zeta) \\ &= \sum_i \int_{-\infty}^{+\infty} s_{ijt} \left[\lambda_j - \sum_l s_{ilt}\,\lambda_l\right] dF(\zeta).\end{aligned} \tag{H3}$$

Continuous interaction ψ

Because ψ enters through $(\Theta_i + \alpha X_{it})\,\omega_{jt}$, which is common to all physicians in a group and constant across groups only through ω_{jt}, the logit derivative is

$$\begin{aligned}\frac{\partial s_{jt}}{\partial \psi} &= \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \psi}\, dF(\zeta) \\ &= \sum_i \int_{-\infty}^{+\infty} s_{ijt}\,(\Theta_i + \alpha X_{it}) \left[\omega_{jt} - \sum_l s_{ilt}\,\omega_{lt}\right] dF(\zeta).\end{aligned} \tag{H4}$$

Let g denote a specific physician group. For each element g of the η vector, we consider:

- if physician $j \in g$

$$\begin{aligned}\frac{\partial s_{jt}}{\partial \eta_j^{\{g\}}} &= \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \eta_j^{\{g\}}}\, dF(\zeta) \\ &= \sum_i \int_{-\infty}^{+\infty} s_{ijt}(\Theta_i + \alpha X_{it}) \left[\tilde{\rho}_{ijt}(1-\tilde{\rho}_{ijt}) - \sum_{l\in g} \tilde{\rho}_{ilt}(1-\tilde{\rho}_{ilt}) s_{ilt}\right] dF(\zeta).\end{aligned} \tag{H5}$$

- if physician $j \notin g$

$$\frac{\partial s_{jt}}{\partial \eta_j^{\{g\}}} = \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \eta_j^{\{g\}}} \, dF(\zeta)$$
$$= -\sum_i \int_{-\infty}^{+\infty} s_{ijt}(\Theta_i + \alpha X_{it}) \left[\sum_{l \in g} \tilde{\rho}_{ilt}(1 - \tilde{\rho}_{ilt}) s_{ilt} \right] \, dF(\zeta). \tag{H6}$$

Similarly, for the constant term, we have

$$\frac{\partial s_{jt}}{\partial cons} = \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial cons} \, dF(\zeta)$$
$$= \sum_i \int_{-\infty}^{+\infty} s_{ijt}(\Theta_i + \alpha X_{it}) \left[\tilde{\rho}_{ijt}(1 - \tilde{\rho}_{ijt}) - \sum_l \tilde{\rho}_{ilt}(1 - \tilde{\rho}_{ilt}) s_{ilt} \right] \, dF(\zeta). \tag{H7}$$

Let k denote a specific patient group. For each element k of the Θ vector, we have:

$$\frac{\partial s_{jt}}{\partial \Theta_i^{\{k\}}} = \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \Theta_i^{\{k\}}} \, dF(\zeta) \tag{H8}$$

with:

$$\frac{\partial s_{ijt}}{\partial \Theta_i^{\{k\}}} = \begin{cases} s_{ijt} \left[(\Theta_i + X_{it}\alpha) \left(\tilde{\rho}_{ijt}(1 - \tilde{\rho}_{ijt}) - \sum_l s_{ilt} \tilde{\rho}_{ilt}(1 - \tilde{\rho}_{ilt}) \right) + \tilde{\rho}_{ijt} - \sum_l s_{ilt} \tilde{\rho}_{ilt} \right] & \text{if } i \in k \\ 0 & \text{if } i \notin k \end{cases}$$

For each element m of α, we have

$$\frac{\partial s_{jt}}{\partial \alpha^{\{m\}}} = \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \alpha^{\{m\}}} \, dF(\zeta)$$
$$= \sum_i \int_{-\infty}^{+\infty} s_{ijt} X_{it}^{\{m\}} \left[(\Theta_i + X_{it}\alpha) \left(\tilde{\rho}_{ijt}(1 - \tilde{\rho}_{ijt}) - \sum_l s_{ilt} \tilde{\rho}_{ilt}(1 - \tilde{\rho}_{ilt}) \right) \right.$$
$$\left. + \tilde{\rho}_{ijt} - \sum_l s_{ilt} \tilde{\rho}_{ilt} \right] \, dF(\zeta). \tag{H9}$$

For the partial derivatives with respect to δ, we have

$$\frac{\partial s_{jt}}{\partial \delta_{jt}} = \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \delta_{jt}} \, dF(\zeta) = \sum_i \int_{-\infty}^{+\infty} s_{ijt}(1 - s_{ijt}) \, dF(\zeta), \tag{H10}$$

and

$$\frac{\partial s_{jt}}{\partial \delta_{lt}} = \sum_i \int_{-\infty}^{+\infty} \frac{\partial s_{ijt}}{\partial \delta_{lt}} \, dF(\zeta) = \sum_i \int_{-\infty}^{+\infty} -s_{ijt} s_{ilt} \, dF(\zeta). \tag{H11}$$

The fixed point imposes that the predicted shares are equal to the observed shares. That is, $0 = s_{jt}(\Xi) - \mathcal{S}_{jt}$. By the implicit function theorem, we have that

$$\left(\frac{d\boldsymbol{\delta}_t}{d\sigma_\zeta} \quad \frac{d\boldsymbol{\delta}_t}{d\sigma_\lambda} \quad \frac{d\boldsymbol{\delta}_t}{d\psi} \quad \frac{d\boldsymbol{\delta}_t}{d\boldsymbol{\phi}} \quad \frac{d\boldsymbol{\delta}_t}{d\boldsymbol{\eta}} \quad \frac{d\boldsymbol{\delta}_t}{dcons} \quad \frac{d\boldsymbol{\delta}_t}{d\boldsymbol{\Theta}} \quad \frac{d\boldsymbol{\delta}_t}{d\boldsymbol{\alpha}}\right) = -\left(\frac{\partial \boldsymbol{s}_t}{\partial \boldsymbol{\delta}_t}\right)^{-1} \left[\frac{\partial \boldsymbol{s}_t}{\partial \sigma_\zeta} \quad \frac{\partial \boldsymbol{s}_t}{\partial \sigma_\lambda} \quad \frac{\partial \boldsymbol{s}_t}{\partial \psi} \quad \frac{\partial \boldsymbol{s}_t}{\partial \boldsymbol{\phi}} \quad \frac{\partial \boldsymbol{s}_t}{\partial \boldsymbol{\eta}} \quad \frac{\partial \boldsymbol{s}_t}{\partial cons} \quad \frac{\partial \boldsymbol{s}_t}{\partial \boldsymbol{\Theta}} \quad \frac{\partial \boldsymbol{s}_t}{\partial \boldsymbol{\alpha}}\right],$$

Now, we can use these to form the components of the Murphy-Topel adjusted standard errors, which consist of the gradients of the loglikelihood with respect to $\Xi_2 = \left(\sigma_\zeta, \phi^{\{m\}}\right)$ and $\Xi_1 = \left(cons, \eta^{\{g\}}, \Theta^{\{k\}}, \alpha^{\{m\}}\right)$. From Equation (18), we have:

$$\frac{\partial \mathcal{L}_2}{\partial \Xi_2} = \sum_i \frac{\partial}{\partial \Xi_2} \log \int \prod_{t \in T_i} \prod_{j \in J_t} (s_{ijt}(\zeta; \boldsymbol{\Xi}_2, \widehat{\boldsymbol{\Xi}}_1))^{y_{ijt}} \, dF(\zeta) \tag{H12}$$

$$\frac{\partial \mathcal{L}_2}{\partial \Xi_1} = \sum_i \frac{\partial}{\partial \Xi_1} \log \int \prod_{t \in T_i} \prod_{j \in J_t} (s_{ijt}(\zeta; \boldsymbol{\Xi}_2, \widehat{\boldsymbol{\Xi}}_1))^{y_{ijt}} \, dF(\zeta) \tag{H13}$$

Note that we can easily calculate the gradients for the outside good choice probabilities because $s_{i0t} = 1 - \sum_l s_{ilt}$. Therefore, the respective gradients are

$$\frac{\partial s_{i0t}}{\partial \gamma} = -\sum_l \frac{\partial s_{ilt}}{\partial \gamma} \tag{H14}$$

where $\gamma \in \{\sigma_\zeta, \sigma_\lambda, \psi, \phi^{\{m\}}, cons, \eta^{\{g\}}, \Theta^{\{k\}}, \alpha^{\{m\}}\}$.

I Technical Details on Counterfactual Simulations

We consider four counterfactual scenarios. Table I1 presents each of the scenarios. When we exclude physicians, we focus on the counterfactual choice probabilities of the affected patients, holding the choices of non-affected patients fixed. We need s_{ijt} and ρ_{ijt} of the affected patients to examine the counterfactual physician choice and opioid prescriptions.

Counterfactual type	**Did s_{ijt} change?**	**Did ρ_{ijt} change?**	**Affected Patients**
Exclude and resort	Yes	No	Patients that chose excluded physicians
Exclude and randomize	Yes	No	Patients that chose excluded physicians
Exclude, resort and break bad sorting	Yes, through $\tilde{\rho}$	Yes	Patients that chose excluded physicians
Break bad sorting	Yes, through $\tilde{\rho}$	Yes	All patients
Patient information, $\sigma_\zeta \neq \hat{\sigma}_\zeta$	Yes	No	All patients

Table I1: Counterfactual simulations

I.1 Exclude top % physician and affected patients resort

We simulate the exclusion policy by dropping physicians in the top n^{th} percentile of risk-adjusted prescription rate. We set the mean utility δ_{jt} of the excluded physicians as $-\infty$. We use Gaussian quadrature to calculate the choice probabilities s_{ijt}. Thus, the patient's perceived prescription probabilities $\tilde{\rho}_{ijtm}$ at each integration node m is:

$$\tilde{\rho}_{ijtm} = \frac{\exp(\eta_j + \Theta_i + \ \tau_t) * \exp(\zeta_{jm})}{1 + \exp(\eta_j + \Theta_i + \ \tau_t) * \exp(\zeta_{jm})}$$

where ζ_{jm} is the integration node. The patient's choice probabilities s_{ijtm} are:

$$s_{ijtm} = \frac{\exp(\delta_{jt} + \tilde{\rho}_{ijtm} \times (\Theta_i + \ \tau_t) + Z_{it}\phi)}{1 + \sum_k \exp(\delta_{kt} + \tilde{\rho}_{iktm} \times (\Theta_i + \ \tau_t) + Z_{it}\phi)}$$

I.2 Exclude top % physician and randomize affected patient choice

For this counterfactual, we use Monte Carlo integration to compute the choice probabilities s_{ijt}. At each Monte Carlo draw m, we have $J+1$ T1EV errors and one standard normal draw for each patient i, scaled by the dispersion of information type, σ_ζ. The perceived prescription probabilities are:

$$\tilde{\rho}_{ijtm} = \frac{\exp(\eta_j + \Theta_i + \ \tau_t) * \exp(\zeta_{jm})}{1 + \exp(\eta_j + \Theta_i + \ \tau_t) * \exp(\zeta_{jm})}$$

The indirect utility of a non-excluded physician is given by:

$$u_{ijtm} = \delta_{jt} + [\tilde{\rho}_{ijtm} * (\Theta_i + \ \tau_t)] + Z_{it}\phi + \nu_{jtm}$$

We calculate the weighted-average utility of all non-excluded physicians and compare it with the realized T1EV error of the outside option. The weights are the conditional market share of the non-excluded physicians. Patients choose the outside option if and only if:

$$\nu_{0tm} \geq \bar{u}_{itm}$$

where ν_{0tm} is the T1EV of the outside option in market t at MC draw m and $\bar{u}_{itm}$ is the weighted average utility of the non-excluded physicians and takes the form:

$$\bar{u}_{itm} = \sum_{j \notin \text{excluded}} s_{jt}^{cond} \times u_{ijtm}$$

s_{jt}^{cond} is the conditional market share of the non-excluded physicians. We implement this 100 times per patient and use the average across the 100 trials to calculate the counterfactual probability of choosing the outside option, s_{i0t}. We calculate the counterfactual choice probabilities s_{ijt} by multiplying $1 - s_{i0t}$ with the conditional market shares of the non-excluded physicians.

I.3 Exclude top % physician, affected patients resort and break bad sorting

This counterfactual exercise is similar to section I.1, where we exclude physicians and allow patients to resort. Additionally, we require physicians to prescribe only when considering preferences that are concordant with CDC guidelines or medical consensus. Specifically, for each patient affected by the exclusion policy, we require physicians to prescribe based on the following rules:

1. If the patient suffers from sickle cell disease, cancer-related pain, or is in end-of-life care or palliative care, the prescription rates are unchanged:

$$\rho_{ijt} = \frac{\exp\left(\eta_j + \overbrace{X_{i1}\beta_1 + X_{i2}\beta_2 + \theta_i}^{\Theta_i} + \tau_t\right)}{1 + \exp\left(\eta_j + \underbrace{X_{i1}\beta_1 + X_{i2}\beta_2 + \theta_i}_{\Theta_i} + \tau_t\right)}$$

2. In all other cases, if a patient's preference related to bad sorting is positive, that is, the sum of observed preferences concordant with neither the guidelines nor medical consensus, plus a portion of θ_i (depending on assumptions about the share unsupported by medical consensus), the physician should ignore this preference. If the patient's preference related to bad sorting is negative, we do not intervene, and the physician should take this preference into account:

$$\rho_{ijt} = \frac{\exp\left(\eta_j + \overbrace{X_{i1}\beta_1 + \alpha\theta_i}^{\Theta_i} + \tau_t\right)}{1 + \exp\left(\eta_j + \underbrace{X_{i1}\beta_1 + \alpha\theta_i}_{\Theta_i} + \tau_t\right)}$$

where $\alpha \in [0, 1]$ and is the portion of θ_i that is supported by medical consensus. As a result, the

choice probabilities s_{ijt} for the second group of patients are:

$$s_{ijt} = \frac{\exp\left(\delta_{jt} + \overbrace{\hat{\hat{\rho}}_{ijt}}^{f(\hat{\Theta}_i = X_{i1}\hat{\beta}_1 + \alpha\theta_i)} \times (\overbrace{\Theta_i}^{unchanged} + \tau_t) + Z_{it}\phi\right)}{1 + \sum_k \exp\left(\delta_{kt} + \underbrace{\hat{\hat{\rho}}_{ikt}}_{f(\hat{\Theta}_i = X_{i1}\hat{\beta}_1 + \alpha\theta_i)} \times (\underbrace{\Theta_i}_{unchanged} + \tau_t) + Z_{it}\phi\right)}$$

I.4 Break bad sorting

Similar to section I.3, we require physicians to ignore patient opioid preferences that are neither concordant with CDC guidelines nor with medical consensus if the patient does not suffer from sickle cell disease, cancer-related pain, or end-of-life care or palliative care. We do not exclude any physicians and allow patients to resort, knowing that physicians are now strictly following the CDC guidelines for opioid prescriptions. We examine the counterfactual choice probabilities of all patients.

I.5 Patient information, $\sigma_\zeta \neq \hat{\sigma}_\zeta$

In this counterfactual, we simulate different information dispersion in patients' perceptions of the physician's prescribing propensity. We do not exclude any physicians and allow patients to resort. We examine the counterfactual choice probabilities of all patients.

I.6 Decomposing the drop in aggregate prescriptions

Table I2: Excluding top 10% of physicians

	Resorting	No Resorting
$\Delta_{baseline}$ (pp)	-6.7	-5.2
– Forego Care	-4.4 (65%)	-4.2 (81%)
– Switch Physicians	-2.4 (35%)	-1.0 (19%)

Note: To quantify foregone care and switch-physician channels, we implement versions of CF1 and CF2 where we normalize the patient's outside option choice probability to equal the baseline level. Denote Y^k_{all} as the outcome of counterfactual k when patients can choose both inside and outside options and Y^k_{in} as the outcome of counterfactual k when patients can only choose the inside options. In the "resorting" column, the switch-physician channel is identified by Y^1_{in}, and the forego care channel is identified by $Y^1_{in} - Y^1_{all}$. In the "no resorting" column, the switch physician channel is identified by $Y^1_{in} - Y^2_{in}$ and the forego care channel is identified by $(Y^1_{in} - Y^1_{all}) - (Y^2_{in} - Y^2_{all})$

J Diversion Ratios and the Physician Substitution Matrix

We summarize the estimated demand system through a *diversion-ratio matrix* (Conlon and Mortimer, 2021): for each physician, we remove them from the choice set one at a time, track where their displaced demand goes, and aggregate the resulting single-physician diversion – weighted by the demand each physician displaced – to the level of physician groups. This section describes Figure **??** and derives the closed-form expression used to compute it.

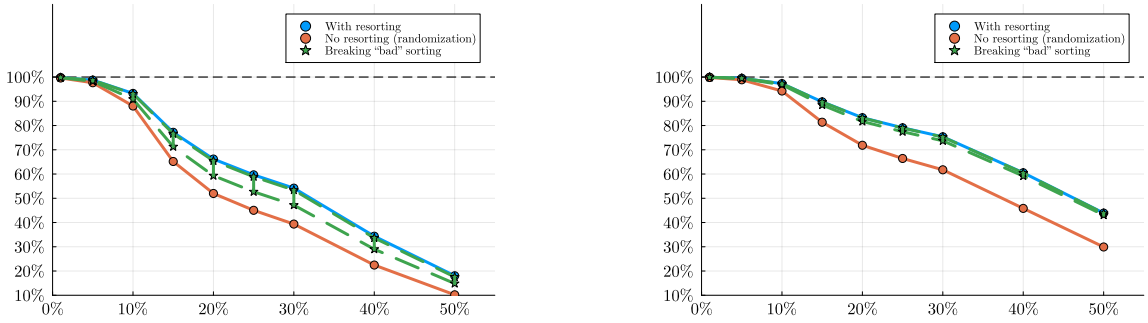
J.1 Second-choice diversion ratio

We compute second-choice diversion ratios following Conlon and Mortimer (2021), removing one physician at a time and tracking where their displaced demand goes. For example, j is the origin physician and k is the destination physician.

$$D_{jk} = \frac{s'_k - s_k}{s_j} = \frac{\int \frac{s_{ij}\, s_{ik}}{1 - s_{ij}}\, \partial F_i}{\int s_{ij}\, \partial F_i} = \frac{\int s'_{ik} - s_{ik}\, \partial F_i}{\int s_{ij}\, \partial F_i}$$

K Counterfactual Policies by Different Exclusion Cutoffs

Figure K1: Counterfactual results by physician exclusion cutoff



(a) Aggregate prescriptions

(b) Aggregate visits

Notes: This figure plots aggregate prescriptions and visits across all three counterfactuals as a function of the exclusion percentile. In the breaking "bad" sorting counterfactual, we require physicians to prescribe based on preferences concordant with CDC guidelines and appropriate medical consensus consistent with assumption A or B. All numbers are expressed as percentages relative to the baseline.